

FortiOS™ Handbook - High Availability

FortiOS 5.4.5



FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

FORTICAST

<http://forticast.fortinet.com>

END USER LICENSE AGREEMENT

<http://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdocs@fortinet.com



Friday, August 25, 2017

FortiOS™ Handbook - High Availability

01-545-99686-20170825

TABLE OF CONTENTS

Change Log	13
Introduction	15
Before you begin	15
Before you set up a cluster	15
How this guide is organized	16
FortiOS 5.4 HA new features	16
High Availability hello-holddown CLI option typo fix (382364)	16
HA diagnose checksum command changes (259710)	16
FGCP supports BFD enabled BGP graceful restart after an HA failover (255574)	17
FRUP is not supported by FortiOS 5.4 (295198)	17
VOIP application control sessions are no longer blocked after an HA failover (273544)	17
Firewall local-in policies are supported for the dedicated HA management interface (276779 246574)	17
HA heartbeat traffic set to the same priority level as data traffic (276665)	17
FGSP CLI command name changed (262340)	18
FGSP now supports synchronizing IPsec sessions (262340)	18
Monitoring VLAN interfaces (220773)	18
Improvements to the get system ha status command output (259416)	18
FortiGate HA cluster support for managed switches (276488 266084)	18
HA cluster health displayed on the Unit Operation dashboard widget (260547)	19
Solving the High Availability problem	20
FortiGate Cluster Protocol (FGCP)	20
FortiGate Session Life Support Protocol (FGSP)	21
VRRP	22
Session-Aware Load Balancing Clustering (SLBC)	23
Enhanced Load Balancing Clustering (ELBC)	24
Content clustering	24
An introduction to the FGCP	26
About the FGCP	26
FGCP failover protection	28
Session Failover	28
Load Balancing	28
Virtual Clustering	28
Full Mesh HA	28

Cluster Management	29
Synchronizing the configuration (and settings that are not synchronized)	30
Preparing the FortiGates before setting up a FGCP cluster	30
DHCP and PPPoE	30
Firmware version	30
About HA and licensing	30
FortiOS Carrier license	31
Support contracts and FortiGuard, FortiCloud, FortiClient, VDOMs Licensing	31
FortiToken Licenses	31
Certificates	32
Configuring FortiGates for FGCP HA operation	32
Connecting a FortiGate HA cluster	34
Verifying the cluster status from the Unit Operation dashboard widget	35
Active-passive and active-active HA	36
Active-passive HA (failover protection)	36
Active-active HA (load balancing and failover protection)	37
Identifying the cluster and cluster units	37
Group name	37
Password	38
Group ID	38
Device failover, link failover, and session failover	38
Primary unit selection	39
Viewing how the primary unit was selected	40
Primary unit selection and monitored interfaces	41
Primary unit selection and age	41
Primary unit selection and device priority	44
Primary unit selection and the FortiGate serial number	46
Points to remember about primary unit selection	46
Temporarily setting a cluster unit to be the primary unit	46
HA override	47
Override and primary unit selection	48
Controlling primary unit selection using device priority and override	49
Points to remember about primary unit selection when override is enabled	49
Configuration changes can be lost if override is enabled	49
Override and disconnecting a unit from a cluster	50
Delaying how quickly the primary unit rejoins the cluster when override is enabled	50
FortiGate HA compatibility with DHCP and PPPoE	50
HA and distributed clustering	51
Clusters of three or four FortiGates	52
Connecting a cluster of three FortiGates	53
Disk storage configuration and HA	56
FGCP high availability best practices	57

Heartbeat interfaces.....	57
Interface monitoring (port monitoring).....	58
FGCP HA terminology.....	59
HA GUI options.....	62
Mode.....	63
Device Priority.....	63
Reserve Management Port for Cluster Member.....	63
Do NOT Synchronize Management VDOM Configuration.....	63
Group Name.....	64
Password.....	64
Enable Session pickup.....	64
Port Monitor.....	64
Heartbeat Interface.....	64
VDOM partitioning.....	65
FGCP configuration examples and troubleshooting.....	66
About the examples in this chapter.....	66
How to set up FGCP clustering (recommended steps).....	66
1. Configuring the primary FortiGate.....	67
2. Configuring the backup FortiGate.....	69
3. Connecting the cluster.....	72
4. Checking cluster operation and disabling override.....	72
5. Results.....	73
Setting up two new FortiGates as an FGCP cluster.....	73
Example NAT/Route mode HA network topology.....	73
General configuration steps.....	74
Configuring a NAT/Route mode active-passive cluster of two FortiGates - GUI.....	75
Configuring a NAT/Route mode active-passive cluster of two FortiGates - CLI.....	79
Adding a new FortiGate to an operating cluster.....	84
Active-active HA cluster in Transparent mode.....	85
Example Transparent mode HA network topology.....	85
General configuration steps.....	86
Configuring a Transparent mode active-active cluster of two FortiGates - GUI.....	87
Configuring a Transparent mode active-active cluster of two FortiGates - CLI.....	91
FortiGate-5000 active-active HA cluster with FortiClient licenses.....	97
Example network topology.....	98
Configuring the FortiGate-5000 active-active cluster - GUI.....	99
Configuring the FortiGate-5000 active-active cluster - CLI.....	104
Converting a standalone FortiGate to a cluster.....	108
1. Adding the backup FortiGate and configuring HA.....	109
2. Results.....	113
Replacing a failed cluster unit.....	113
FGCP HA with 802.3ad aggregated interfaces.....	115

HA interface monitoring, link failover, and 802.3ad aggregation	115
HA MAC addresses and 802.3ad aggregation	116
Link aggregation, HA failover performance, and HA mode	116
General configuration steps	117
Configuring active-passive HA cluster that includes aggregated interfaces - GUI	117
Configuring active-passive HA cluster that includes aggregate interfaces - CLI	122
Example HA and redundant interfaces	126
HA interface monitoring, link failover, and redundant interfaces	127
HA MAC addresses and redundant interfaces	127
Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode	127
Connecting multiple redundant interfaces to one switch while operating in active-active HA mode	128
General configuration steps	128
Configuring active-passive HA cluster that includes redundant interfaces - GUI	128
Configuring active-passive HA cluster that includes redundant interfaces - CLI	133
Troubleshooting HA clusters	137
Ignoring hardware revisions	137
Before you set up a cluster	138
Troubleshooting the initial cluster configuration	139
More troubleshooting information	141
Virtual clusters	143
Virtual clustering overview	143
Virtual clustering and failover protection	143
Virtual clustering and heartbeat interfaces	143
Virtual clustering and HA override	144
Virtual clustering and load balancing or VDOM partitioning	144
Configuring HA for virtual clustering	145
Example virtual clustering with two VDOMs and VDOM partitioning	147
Example virtual clustering network topology	147
General configuration steps	148
Configuring virtual clustering with two VDOMs and VDOM partitioning - GUI	149
Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI	154
Example inter-VDOM links in a virtual clustering configuration	159
Configuring inter-VDOM links in a virtual clustering configuration	160
Troubleshooting virtual clustering	162
Full mesh HA	163
Full mesh HA overview	163
Full mesh HA and redundant heartbeat interfaces	164
Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces	164
Example full mesh HA configuration	165
Full mesh HA configuration	166

Full mesh switch configuration.....	166
Full mesh network connections.....	166
How packets travel from the internal network through the full mesh cluster and to the Internet.....	166
Configuring full-mesh HA - GUI.....	167
Configuring Full Mesh HA - CLI.....	172
Troubleshooting full mesh HA.....	176
Operating clusters and virtual clusters.....	177
Operating a cluster.....	177
Operating a virtual cluster.....	177
Managing individual cluster units using a reserved management interface.....	178
Using the HA reserved management interface for FortiSandbox, SNMP and other management services.....	179
Configuring the reserved management interface and SNMP remote management of individual cluster units.....	179
Adding firewall local-in policies for the dedicated HA management interface.....	183
Managing individual cluster units in a virtual cluster.....	183
The primary unit acts as a router for subordinate unit management traffic.....	183
Cluster communication with RADIUS and LDAP servers.....	184
Clusters and FortiGuard services.....	184
FortiGuard and active-passive clusters.....	185
FortiGuard and active-active clusters.....	185
FortiGuard and virtual clustering.....	185
Clusters and logging.....	185
Viewing and managing log messages for individual cluster units.....	186
HA log messages.....	186
FortiGate HA message "HA master heartbeat interface <intf_name> lost neighbor information".....	187
Formatting cluster unit hard disks (log disks).....	188
Clusters and SNMP.....	188
SNMP get command syntax for the primary unit.....	189
SNMP get command syntax for any cluster unit.....	190
Getting serial numbers of cluster units.....	191
SNMP get command syntax - reserved management interface enabled.....	191
Adding FortiClient licenses to a cluster.....	192
Adding FortiClient licenses to cluster units with a reserved management interface.....	192
Adding FortiClient licenses to cluster units with no reserved management interface.....	192
Viewing FortiClient license status and active FortiClient users for each cluster unit.....	193
Cluster members list.....	194
Virtual cluster members list.....	194
Viewing HA statistics.....	194
Changing the HA configuration of an operating cluster.....	195
Changing the HA configuration of an operating virtual cluster.....	195

Changing the subordinate unit host name and device priority.....	196
Upgrading cluster firmware.....	196
Changing how the cluster processes firmware upgrades.....	197
Synchronizing the firmware build running on a new cluster unit.....	197
Downgrading cluster firmware.....	198
Backing up and restoring the cluster configuration.....	199
Monitoring cluster units for failover.....	199
Viewing cluster status from the CLI.....	199
Get system ha status example - two FortiGates in active-passive mode.....	203
Get system ha status example - three FortiGates in active-active mode.....	204
Get system ha status example - virtual cluster.....	205
About the HA cluster index and the execute ha manage command.....	208
Managing individual cluster units.....	210
Disconnecting a cluster unit from a cluster.....	211
Adding a disconnected FortiGate back to its cluster.....	212
HA diagnose commands.....	213
all-xdb.....	213
all-vcluster.....	214
stat.....	214
HA and failover protection.....	215
About active-passive failover.....	215
Device failure.....	215
Link failure.....	215
Primary unit recovery.....	216
About active-active failover.....	216
Device failover.....	216
HA heartbeat and communication between cluster units.....	217
Heartbeat interfaces.....	218
Connecting HA heartbeat interfaces.....	219
Heartbeat packets and heartbeat interface selection.....	219
Interface index and display order.....	220
HA heartbeat interface IP addresses.....	220
Heartbeat packet Ethertypes.....	221
Modifying heartbeat timing.....	222
Enabling or disabling HA heartbeat encryption and authentication.....	223
Heartbeat bandwidth requirements.....	224
Cluster virtual MAC addresses.....	224
Changing how the primary unit sends gratuitous ARP packets after a failover.....	225
Disabling gratuitous ARP packets after a failover.....	226
How the virtual MAC address is determined.....	226
Displaying the virtual MAC address.....	228
Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain.....	229

Synchronizing the configuration.....	232
Disabling automatic configuration synchronization.....	232
Incremental synchronization.....	233
Periodic synchronization.....	233
Console messages when configuration synchronization succeeds.....	234
Console messages when configuration synchronization fails.....	234
Comparing checksums of cluster units.....	236
How to diagnose HA out of sync messages.....	238
Recalculating the checksums to resolve out of sync messages.....	239
Synchronizing kernel routing tables.....	239
Controlling how the FGCP synchronizes kernel routing table updates.....	240
Configuring graceful restart for dynamic routing failover.....	241
Graceful OSPF restart.....	242
Graceful BGP restart.....	242
Notifying BGP neighbors when graceful restart is enabled.....	242
Bidirectional Forwarding Detection (BFD) enabled BGP graceful restart.....	243
Link failover (port monitoring or interface monitoring).....	243
If a monitored interface on the primary unit fails.....	245
If a monitored interface on a subordinate unit fails.....	245
How link failover maintains traffic flow.....	245
Recovery after a link failover and controlling primary unit selection (controlling falling back to the prior primary unit).....	246
Preventing a primary unit change after a failed link is restored.....	247
Testing link failover.....	247
Updating MAC forwarding tables when a link failover occurs.....	247
Multiple link failures.....	248
Example link failover scenarios.....	248
Monitoring VLAN interfaces.....	249
Subsecond failover.....	249
Remote link failover.....	250
Adding HA remote IP monitoring to multiple interfaces.....	252
Changing the link monitor failover threshold.....	253
Monitoring multiple IP addresses from one interface.....	254
Flip timeout.....	254
Restoring normal cluster operation after the remote link is restored.....	254
Detecting HA remote IP monitoring failovers.....	255
Failover and attached network equipment.....	255
Monitoring cluster units for failover.....	255
NAT/Route mode active-passive cluster packet flow.....	255
Packet flow from client to web server.....	256
Packet flow from web server to client.....	257
When a failover occurs.....	257

Transparent mode active-passive cluster packet flow.....	257
Packet flow from client to mail server.....	258
Packet flow from mail server to client.....	259
When a failover occurs.....	259
Failover performance.....	260
Device failover performance.....	260
Link failover performance.....	260
Reducing failover times.....	260
Session failover (session-pickup).....	262
Enabling session-pickup for TCP, UDP and ICMP session failover.....	262
If session pickup is disabled.....	263
Improving session synchronization performance.....	263
Reducing the number of sessions that are synchronized.....	263
Using multiple FortiGate interfaces for session synchronization.....	264
Session failover limitations for sessions passing through the cluster.....	264
More about TCP session failover.....	265
UDP, ICMP, multicast and broadcast packet session failover.....	266
SIP session failover.....	266
FortiOS Carrier GTP session failover.....	266
SSL offloading and HTTP multiplexing session failover.....	267
Active-active HA subordinate units sessions can resume after a failover.....	267
Session failover limitations for sessions terminated by the cluster.....	267
Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and Web Caching session failover.....	268
SSL VPN session failover and SSL VPN authentication failover.....	269
PPTP and L2TP VPN sessions.....	269
Synchronizing IPsec VPN SAs.....	269
Synchronizing SAs for IKEv1.....	269
Synchronizing SAs for IKEv2.....	270
WAN optimization and HA.....	270
HA and load balancing.....	271
Load balancing overview.....	271
Load balancing schedules.....	272
More about active-active failover.....	273
HTTPS sessions, active-active load balancing, and proxy servers.....	273
Selecting a load balancing schedule.....	273
Load balancing TCP and UDP sessions.....	274
Using NP4 or NP6 processors to offload load balancing.....	274
Configuring weighted-round-robin weights.....	275
Dynamically optimizing weighted load balancing according to how busy cluster units are.....	276
Example weighted load balancing configuration.....	278
NAT/Route mode active-active cluster packet flow.....	279

Packet flow from client to web server.....	280
Packet flow from web server to client.....	281
When a failover occurs.....	282
Transparent mode active-active cluster packet flow.....	282
Packet flow from client to mail server.....	283
Packet flow from mail server to client.....	284
When a failover occurs.....	285
HA with FortiGate-VM and third-party products.....	286
FortiGate-VM for VMware HA configuration.....	286
FortiGate VM for Hyper-V HA configuration.....	286
Troubleshooting layer-2 switches.....	287
Forwarding delay on layer 2 switches.....	287
Failover issues with layer-3 switches.....	287
Failover and attached network equipment.....	288
Ethertype conflicts with third-party switches.....	288
LACP, 802.3ad aggregation and third-party switches.....	288
VRRP.....	290
Adding a VRRP virtual router to a FortiGate interface.....	291
VRRP virtual MAC address.....	292
VRRP Groups.....	292
Using a Second Destination IP (VRDST).....	293
Configuring VRRP.....	293
Example VRRP configuration: two FortiGates in a VRRP group.....	293
Example VRRP configuration: VRRP load balancing two FortiGates and two VRRP groups.....	295
Optional VRRP configuration settings.....	296
FortiGate Session Life Support Protocol (FGSP).....	298
Configuring FGSP HA cluster-sync instances.....	300
FGSP clusters with three or more FortiGates.....	300
Selecting the sessions to synchronize.....	301
Synchronizing sessions.....	301
Synchronizing the configuration.....	301
IPsec tunnel synchronization.....	302
Synchronizing UDP and ICMP (connectionless) sessions.....	302
Synchronizing NAT sessions.....	302
Synchronizing asymmetric sessions.....	303
Synchronizing expectation sessions.....	303
Security profile flow-based inspection and asymmetric traffic.....	304
Notes and limitations.....	304
Configuring session synchronization links.....	305
Basic example configuration.....	305
Verifying FGSP configuration and synchronization.....	308

FGSP configuration summary and status.....	308
Verifying that sessions are synchronized.....	309

Change Log

Date	Change Description
25 August, 2017	Corrections to Backing up and restoring the cluster configuration on page 199 .
7 April, 2017	<p>Throughout the document, added information about how all of the FortiGates in a cluster must have the same level of licensing. For example, see About HA and licensing on page 30.</p> <p>Moved and re-organized all of the session failover/session-pickup information into a new chapter, see Session failover (session-pickup) on page 262.</p> <p>Changes to the chapter Virtual clusters on page 143. A virtual clustering configuration can include more than two FortiGates. If a virtual cluster configuration includes a VLAN that is in a different VDOM than the physical interface that the VLAN has been added to, both of these VDOMs must be in the same virtual cluster (vcluster 1 or vcluster 2).</p>
6 February, 2017	New section Using the HA reserved management interface for FortiSandbox, SNMP and other management services on page 179 . Changes to HA and distributed clustering on page 51 , FGCP high availability best practices on page 57 , Heartbeat bandwidth requirements on page 224 , Session failover (session-pickup) on page 262 , and Session failover (session-pickup) on page 262 .
9 January, 2017	<p>A new FortiOS 5.4 feature we missed is that the <code>diagnose sys ha showcsum</code> command changed to <code>diagnose sys ha checksum</code>. As well <code>diagnose sys ha showcsum recalculate</code> changed to <code>diagnose sys ha checksum recalculate</code>. Added a note about this to FortiOS 5.4 HA new features on page 16 and updated the section How to diagnose HA out of sync messages on page 238 and Recalculating the checksums to resolve out of sync messages on page 239.</p> <p>Improved the description of how proxy-based and flow-based security profiles affect session failover in Session failover (session-pickup) on page 262.</p> <p>Included a step for using the <code>hbdev HA</code> option to configure the synchronization link when configuring FGSP in FortiGate Session Life Support Protocol (FGSP) on page 298.</p>
6 December, 2016	<p>New section Restoring normal cluster operation after the remote link is restored on page 254.</p> <p>Multiple changes to the FortiGate Session Life Support Protocol (FGSP) on page 298 to reflect that FGSP supports clusters of up to 16 FortiGates.</p>
2 November 2016	Updated for FortiOS 5.4.2. Changed the spelling of the <code>hello-holddown</code> CLI option to <code>hello-holddown</code> throughout.

Date	Change Description
13 October, 2016	<p>New section Viewing how the primary unit was selected on page 40.</p> <p>Updates related to changes to the <code>get system ha status</code> command throughout the document, especially in Viewing cluster status from the CLI on page 199.</p> <p>More improvements for information about FortiToken licensing and HA clusters including the new section FortiToken Licenses on page 31.</p> <p>Fixes to links and cross references throughout the document.</p>
7 October, 2016	<p>The section Synchronizing asymmetric sessions on page 303 was renamed and updated. New section added: Synchronizing expectation sessions on page 303.</p> <p>More information added to Ignoring hardware revisions on page 137. New section Delaying how quickly the primary unit rejoins the cluster when override is enabled on page 50. Changes to Synchronizing kernel routing tables on page 239 and Configuring graceful restart for dynamic routing failover on page 241.</p>
19 July, 2016	Corrected information about FortiToken licensing and HA clusters throughout.
8 June, 2016	<p>The new section FortiGate Session Life Support Protocol (FGSP) on page 298 contains more details about FGSP IPsec tunnel synchronization. Corrected the information about using the <code>execute ha manage</code> command in Operating clusters and virtual clusters on page 177.</p>
11 February, 2016	Changes to HA and failover protection on page 215 . Corrections to Full mesh HA on page 163 to add that full mesh HA requires 802.1Q (Dot1Q) or ISL communication between redundant switches.
December 15, 2015	New FortiOS 5.4.0 release.

Introduction

This document describes FortiGate HA, the FortiGate Clustering Protocol (FGCP), The FortiGate Session Life Support Protocol (FGSP) and FortiGate VRRP.

Before you begin

Before you begin using this guide, take a moment to note the following:

- If you enable virtual domains (VDOMs), HA is configured globally for the entire FortiGate and the configuration is called virtual clustering.
- This HA guide is based on the assumption that you are a FortiGate administrator. It is not intended for others who may also use the FortiGate, such as FortiClient administrators or end users.
- The configuration examples show steps for both the GUI (GUI) and the CLI.

At this stage, the following installation and configuration conditions are assumed:

- You have two or more FortiGates of the same model available for configuring and connecting to form an HA cluster. You have a copy of the QuickStart Guide for the FortiGates.
- You have administrative access to the GUI and CLI.

Many of the configuration examples in this document begin FortiGates unit configured with the factory default configuration. This is optional, but may make the examples easier to follow. As well, you do not need to have installed the FortiGates on your network before using the examples in this document.

Before you set up a cluster

Before you set up a cluster ask yourself the following questions about the FortiGates that you are planning to use to create a cluster.

Do all the FortiGates have the same hardware configuration? Including the same hard disk configuration and the same optional components installed in the same slots?

1. Do all FortiGates have the same firmware build?
2. Are all FortiGates set to the same operating mode (NAT or Transparent)?
3. Are all the FortiGates operating in the same VDOM mode?
4. If the FortiGates are operating in multiple VDOM mode do they all have the same VDOM configuration?



In some cases you may be able to form a cluster if different FortiGates have different firmware builds, different VDOM configurations, and are in different operating modes. However, if you encounter problems they may be resolved by installing the same firmware build on each unit, and give them the same VDOM configuration and operating mode.

How this guide is organized

This document contains detailed information about how FortiGate HA and the FortiGate Clustering Protocol (FGCP) works. This document all describes all FortiGate HA configuration options, contains detailed configuration examples, and describes how to operate FortiGate clusters. Future versions of this document will include more and more configuration examples and more information about HA functionality.

This FortiOS Handbook chapter contains the following sections:

[Solving the High Availability problem](#) describes the high availability problem and introduces the FortiOS solutions described in this document (FGCP, VRRP, and standalone session synchronization).

[An introduction to the FGCP](#) introduces the FGCP clustering protocol and many of its features and terminology.

[FGCP configuration examples and troubleshooting](#) describes configuring HA clusters and contains HA clustering configuration examples.

[Virtual clusters](#) describes configuring HA virtual clusters and contains virtual clustering configuration examples.

[Full mesh HA](#) describes configuring FortiGate Full mesh HA and contains a full mesh HA configuration example.

[Operating clusters and virtual clusters](#) describes how to operate a cluster and includes detailed information about how various FortiGate systems operate differently in a cluster.

[HA and failover protection](#) describes in detail how FortiGate HA device failover, link failover, and session failover work.

[HA and load balancing](#) describes how FGCP HA active-active load balancing works and how to configure it.

[HA with FortiGate-VM and third-party products](#) describes how FortiGates interact with third-party products.

[VRRP](#) describes FortiOS support of the Virtual Router Redundancy Protocol (VRRP) and its use for high availability.

[FortiGate Session Life Support Protocol \(FGSP\)](#) describes how to use the FGSP feature to support using external routers or load balancers to distribute or load balance sessions between two peer FortiGates.

FortiOS 5.4 HA new features

High Availability hello-holddown CLI option typo fix (382364)

The `config system ha` option `helo-holddown` has been renamed `hello-holddown`.

HA diagnose checksum command changes (259710)

The following diagnose commands changed:

- `diagnose sys ha showcsum-->diagnose sys ha checksum show`
- `diagnose sys ha csum-recalculate-->diagnose sys ha checksum recalculate`
- `diagnose sys ha cached-csum-->diagnose sys ha checksum cached`
- `diagnose sys ha cluster-csum-->diagnose sys ha checksum cluster`

FGCP supports BFD enabled BGP graceful restart after an HA failover (255574)

If an HA cluster is part of a Border Gateway Protocol (BGP) bidirectional forwarding detection (BFD) configuration where both the cluster and the BGP static neighbor are configured for graceful restart, after an HA failover BGP enters graceful restart mode and both the cluster and the BGP neighbor keep their BGP routes.

To support HA and BFD enabled BGP graceful:

- From the cluster, you can add a BFD enabled BGP neighbor as a static BFD neighbor using the `config router bfd` command. Set the BGP auto-start timer to 5 seconds so that after an HA failover BGP on the new primary unit waits for 5 seconds before connect to its BFD neighbors, and then registers BFD requests after establishing the connections. With static BFD neighbors, BFD requests and sessions can be created as soon as possible after the failover. The command `get router info bfd requests` shows the BFD peer requests.
- The BFD session created for a static BFD neighbor/peer request initializes its state as INIT instead of DOWN and its detection time as `bfd-required-min-rx * bfd-detect-mult msec`s.
- When a BFD control packet with a nonzero Your Discriminator (`your_discr`) value is received, if no session can be found to match the `your_discr`, instead of discarding the packet, other fields in the packet, such as addressing information, are used to choose one session that was just initialized, with zero as its remote discriminator.
- When a BFD session in the UP state receives a control packet with zero as Your Discriminator and DOWN as State, the session changes its state to DOWN but will not notify this DOWN event to BGP and/or other registered clients.

FRUP is not supported by FortiOS 5.4 (295198)

With the changes to switch mode, FRUP is no longer available on the FortiGate-100D.

VOIP application control sessions are no longer blocked after an HA failover (273544)

After an HA failover, VoIP sessions that are being scanned by application control will now continue with only a minor interruption, if any. To support this feature, IPS UDP expectation tables are now synchronized between cluster units.

Firewall local-in policies are supported for the dedicated HA management interface (276779 246574)

To add local in polices for the dedicated management interface, enable `ha-mgmt-intf-only` and set `intf` to `any`. Enabling `ha-mgmt-intf-only` means the local-in policy applies only to the VDOM that contains the dedicated HA management interface.

```
config firewall local-in-policy
  edit 0
    set ha-mgmt-intf-only enable
    set intf any
    etc...
end
```

HA heartbeat traffic set to the same priority level as data traffic (276665)

Local out traffic, including HA heartbeat traffic, is now set to high priority to make sure it is processed at the same priority level as data traffic. This change has been made because HA heartbeat traffic can be processed by NP6

processors that are also processing data traffic. When HA heartbeat traffic was set to a lower priority it may have been delayed or dropped by very busy NP6 processors resulting in HA failovers.

FGSP CLI command name changed (262340)

The FortiOS 5.2 command `config system session-sync` has been changed in FortiOS 5.4 to `config system cluster-sync`. Otherwise the command syntax is the same and the `config system ha` commands used for FGSP settings have not changed.

FGSP now supports synchronizing IPsec sessions (262340)

The FGSP now synchronizes IPsec tunnels between FortiGates in an FGSP configuration. IPsec tunnel synchronization synchronizes keys and other run time data between the FortiGates in an FGSP configuration. No additional configuration is required to synchronize IPsec sessions. Also you cannot disable IPsec tunnel synchronization.

The FGSP synchronizes IPsec keys and other runtime data but not actual tunnel sessions. This means that if one of the cluster units goes down the cluster unit that is still operating can quickly get IPsec tunnels re-established without re-negotiating them but all existing tunnel sessions on the failed FortiGate have to be restarted on the still operating FortiGate.

IPsec tunnel sync only supports dialup IPsec. The interfaces on both FortiGates that are tunnel endpoints must have the same IP addresses and external routers must be configured to load balance IPsec tunnel sessions to the FortiGates in the cluster.

Monitoring VLAN interfaces (220773)

When operating in HA mode and if you have added VLAN interfaces to the FortiGates in the cluster, you can use the following command to monitor all VLAN interfaces and send a message if one of the VLAN interfaces is found to be down.

```
config system ha-monitor
  set monitor-vlan enable/disable
  set vlan-hb-interval <interval_seconds>
  set vlan-hb-lost-threshold <vlan-lost-heartbeat-threshold>
end
```

Once configured, this feature works by verifying that the primary unit can connect to the subordinate unit over each VLAN. This verifies that the switch that the VLAN interfaces are connected to is configured correctly for each VLAN. If the primary unit cannot connect to the subordinate unit over one of the configured VLANs the primary unit writes a link monitor log message indicating that the named VLAN went down (log message id 20099).

Improvements to the get system ha status command output (259416)

The `get system ha status` command now displays more information about cluster status including HA health status, cluster uptime, how the primary unit was selected and when this happened, override status, HA heartbeat interface activity, and CPU and memory usage for each cluster unit.


FortiGate HA cluster support for managed switches (276488 266084)

Added the capability to support managed switches from a FortiGate HA cluster. If a standby FortiGate becomes active, it automatically establishes connectivity with the managed switches.

HA cluster health displayed on the Unit Operation dashboard widget (260547)

The Unit Operation dashboard widget now includes the serial number and hostname of all of the FortiGate units in the cluster as well as an indication of the sync status of each cluster member.

Unit Operation



HA Cluster Members	Sync Status	Role
FG-5KD3914800344/FG-5KD3914800344	✓	Master
FG-5KD3914800284/FG-5KD3914800284	✓	Slave

Solving the High Availability problem

The basic high availability (HA) problem for TCP/IP networks and security gateways is keeping network traffic flowing. Uninterrupted traffic flow is a critical component for online systems and media because critical business processes quickly come to a halt when the network is down.

The security gateway is a crucial component of most networks since all traffic passes through it. A standalone network security gateway is a single point of failure that is vulnerable to any number of software or hardware problems that could compromise the device and bring all traffic on the network to a halt.

A common solution to the high availability problem is to eliminate the security gateway as single point of failure by introducing redundancy. With two or more redundant security gateways, if one fails, the remaining one or more gateways keep the traffic flowing. FortiOS provides six redundancy solutions: industry standard VRRP as well as five proprietary solutions: FortiGate Cluster Protocol (FGCP) high availability, FortiGate Session Life Support Protocol (FGSP) high availability, Session-Aware Load Balancing Clustering (SLBC), Enhanced Load Balanced Clustering (ELBC) and Content Clustering.

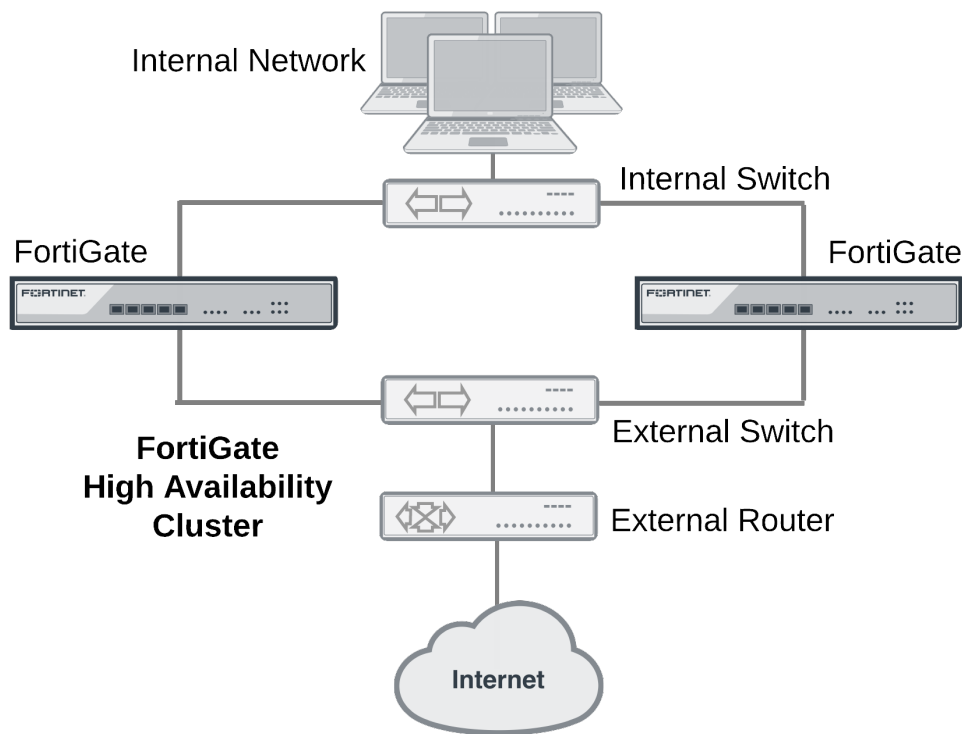


You can combine more than one high availability solution into a single configuration. A common reason for doing this could be to add VRRP to an FGCP or FGSP configuration.

A strong and flexible High availability solution is required for many mission-critical firewall and security profile applications. Each FortiOS high availability solution can be fine tuned to fit into many different network scenarios.

FortiGate Cluster Protocol (FGCP)

FGCP HA provides a solution for two key requirements of critical enterprise networking components: enhanced reliability and increased performance. Enhanced reliability is achieved through device failover protection, link failover protection and remote link failover protection. Also contributing to enhanced reliability is session failover protection for most IPv4 and IPv6 sessions including TCP, UDP, ICMP, IPsec VPN, and NAT sessions. Increased performance is achieved through active-active HA load balancing. Extended FGCP features include full mesh HA and virtual clustering. You can also fine tune the performance of the FGCP to change how a cluster forms and shares information among cluster units and how the cluster responds to failures.

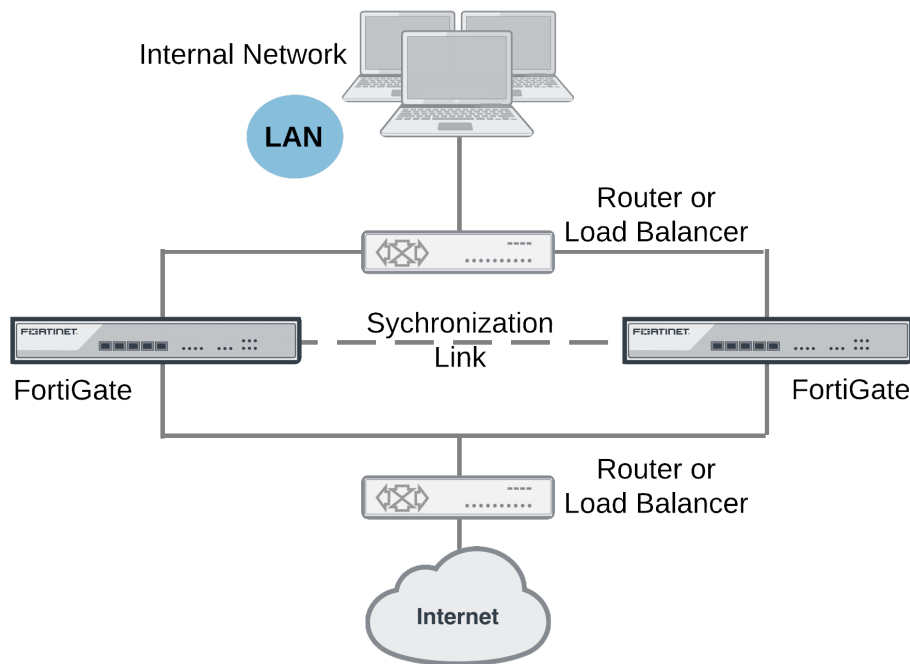


When configured onto your network an FGCP cluster appears to be a single FortiGate operating in NAT/Route or Transparent mode and configuration synchronization allows you to configure a cluster in the same way as a standalone FortiGate. If a failover occurs, the cluster recovers quickly and automatically and also sends administrator notifications so that the problem that caused the failure can be corrected and any failed equipment restored.

The FGCP is compatible with most network environments and most networking equipment. While initial configuration is relatively quick and easy, a large number of tools and configuration options are available to fine tune the cluster for most situations.

FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two identical FortiGates can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute sessions among the FortiGates and the FGSP performs session synchronization of IPv4 and IPv6 TCP, UDP, ICMP, expectation, and NAT sessions to keep the session tables of both FortiGates synchronized.



If one of the FortiGates fails, session failover occurs and active sessions fail over to the unit that is still operating. This failover occurs without any loss of data. As well, the external load balancers or routers detect the failover and re-distribute all sessions to the unit that is still operating.

Load balancing and session failover is done by external routers or load balancers and not by the FGSP. The FortiGates just perform session synchronization which allows session failover to occur without packet loss.

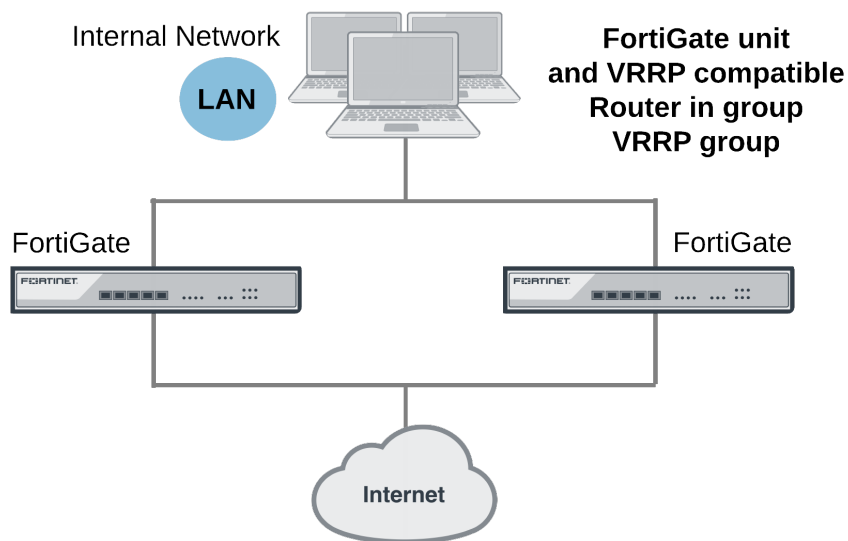
The FGSP also includes configuration synchronization, allowing you to make configuration changes once for both FortiGates instead of requiring duplicate configuration changes on each unit. Settings that identify the FortiGate to the network, for example, interface IP addresses and BGP neighbor settings, are not synchronized so each FortiGate maintains its identity on the network. These settings must be configured separately for each FortiGate.



In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. However, the FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions.

VRRP

FortiGates can function as master or backup Virtual Router Redundancy Protocol (VRRP) routers and can be quickly and easily integrated into a network that has already deployed VRRP. A FortiGate can be integrated into a VRRP group with any third-party VRRP devices and VRRP can provide redundancy between multiple FortiGates.



In a VRRP configuration, when a FortiGate operating as the master unit fails, a backup unit takes its place and continues processing network traffic. If the backup unit is a FortiGate, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate is back on line. You can include different FortiGate models in the same VRRP group.

FortiOS supports VRRP between two or more FortiGates and between FortiGates and third-party routers that support VRRP. Using VRRP you can assign VRRP routers as master or backup routers. The master router processes traffic and the backup routers monitor the master router and can begin forwarding traffic if the master fails. Similar to the FGCP you can configuration VRRP between multiple FortiGates to provide redundancy. You can also create a VRRP group with a FortiGates and any routers that support VRRP.

In a VRRP configuration that consists of one FortiGate and one router, normally the FortiGate would be the master and all traffic would be processed by the FortiGate. If the FortiGate fails, all traffic switches to the router. Network connectivity is maintained even though FortiGate security features will be unavailable until the FortiGate can is back on line.

Session-Aware Load Balancing Clustering (SLBC)

Session-Aware Load Balancing Clusters consist of one or more FortiControllers acting as load balancers and FortiGate-5000s and operating as workers all installed in one or two FortiGate-5000 series chassis.

SLBC clusters load balance TCP and UDP sessions. As a session-aware load balancer, the FortiController includes FortiASIC DP processors that maintain state information for all TCP and UDP sessions. The FortiASIC DP processors are capable of directing any TCP or UDP session to any worker installed in the same chassis. This session-awareness means that all TCP and UDP traffic being processed by a specific worker continues to be processed by the same worker. Session-awareness also means that more complex networking features such as network address translation (NAT), fragmented packets, complex UDP protocols, and complex protocols such as SIP that use pinholes, can be load balanced by the cluster.

For more information about SLBC see the *FortiController Session-Aware Load Balancing Guide*.



You cannot mix FGCP and SLBC clusters in the same FortiGate-5000 chassis.

Enhanced Load Balancing Clustering (ELBC)

ELBC uses FortiSwitch-5000 series load balancers to load balance traffic to FortiGate-5000 workers installed in a FortiGate-5000 chassis. ELBC enhances scalability, reliability, and availability of mission critical IP-based services, such as firewall, antivirus, web filtering, IPS, and so on. It also provides high availability by detecting worker failures and automatically redistributing traffic to the workers that are still operating.

ELBC applies a load balancing algorithm against the source and/or destination address of packets to generate a hash key value. Each worker has hash key values assigned to it. If the workers are running, then the traffic is forwarded to the worker assigned to the hash key. The hash key value generated by the algorithm, the hash keys accepted by the worker blades, and the blade the traffic is sent to are automatically calculated by the FortiSwitch.

For more information about ELBC see the *ELBC Configuration Guide*.



You cannot mix FGCP and ELBC clusters in the same FortiGate-5000 chassis.

Content clustering

A content cluster employs FortiSwitch-5203Bs or FortiController-5902Ds to load balance content sessions to FortiGate-5000 workers. FortiSwitch-5203B content clusters consist of one or more FortiSwitch-5203Bs and multiple FortiGate-5001Bs workers. FortiController-5902D content clusters consist of one or more FortiController-5902Ds and multiple FortiGate-5001B workers.

Operating as a FortiGate in content cluster mode, a primary FortiSwitch-5203B or FortiController-5902D performs routing, firewalling, stateful inspection, IPsec and SSL VPN encryption/decryption, and other FortiGate functions. The FortiSwitch-5203B includes NP4 processors and the FortiController-5902Ds includes NP6 processors and an integrated switch fabrics that provides fastpath acceleration by offloading communication sessions from the FortiGate CPU.

Using content cluster weighted load balancing, the FortiSwitch-5203Bs or FortiController-5902Ds distribute sessions that require content processing to the workers over the FortiGate-5000 chassis fabric backplane. Content processing sessions include proxy and flow-based security profile functions such as virus scanning, intrusion protection, application control, IPS, web filtering, email filtering, and VoIP. Load balancing is offloaded to the NP4 or NP6 processors resulting in improved load balancing performance. In some networks, the NP4 or NP6 processors also allow you to configure the efficiently load balance TCP and UDP sessions.

Content cluster mode is similar to active-active HA where the FortiSwitch-5203B or FortiController-5902D operates as the primary unit and load balances security profile sessions to the workers installed in the chassis using weighted load balancing. In this configuration, the HA mode is active-active, the HA load balancing schedule is weight-round-robin and load-balance-all is disabled. You can adjust the HA weighted load balancing weights to change how sessions are load balanced.

You can add a second FortiSwitch-5203B or FortiController-5902D to a content cluster as a backup. The primary FortiSwitch-5203B or FortiController-5902D can load balance sessions to the backup FortiSwitch-5203B or FortiController-5902D as well as the workers. You can control how many sessions are processed by the backup FortiSwitch-5203B or FortiController-5902D by configuring the HA load balancing weights. You can also configure the content cluster to operate the backup FortiSwitch-5203B or FortiController-5902D in standby mode. In this mode the backup FortiSwitch-5203B or FortiController-5902D does not process any sessions but is just there to take over content clustering if the primary unit fails.

Once the content cluster has been established and all FortiControllers and workers have joined the cluster, you can configure the cluster from the FortiSwitch-5203B or FortiController-5902D GUI or CLI. All configuration changes made to the primary unit are automatically synchronized to all cluster units.

FortiSwitch-5203B or FortiController-5902D firmware upgrades are done from the primary FortiSwitch-5203B or FortiController-5902D GUI or CLI. Worker firmware upgrades are done from the FortiSwitch-5203B or FortiController-5902D CLI where a single firmware image is uploaded once and synchronized to all of the workers.

An introduction to the FGCP

A FortiGate HA cluster consists of two to four FortiGates configured for HA operation. Each FortiGate in a cluster is called a cluster unit. All cluster units must be the same FortiGate model with the same FortiOS firmware build installed. All cluster units must also have the same hardware configuration (for example, the same AMC modules installed in the same slots, the same number of hard disks and so on) and be running in the same operating mode (NAT/Route mode or Transparent mode).



You can create an FGCP cluster of up to four FortiGates.

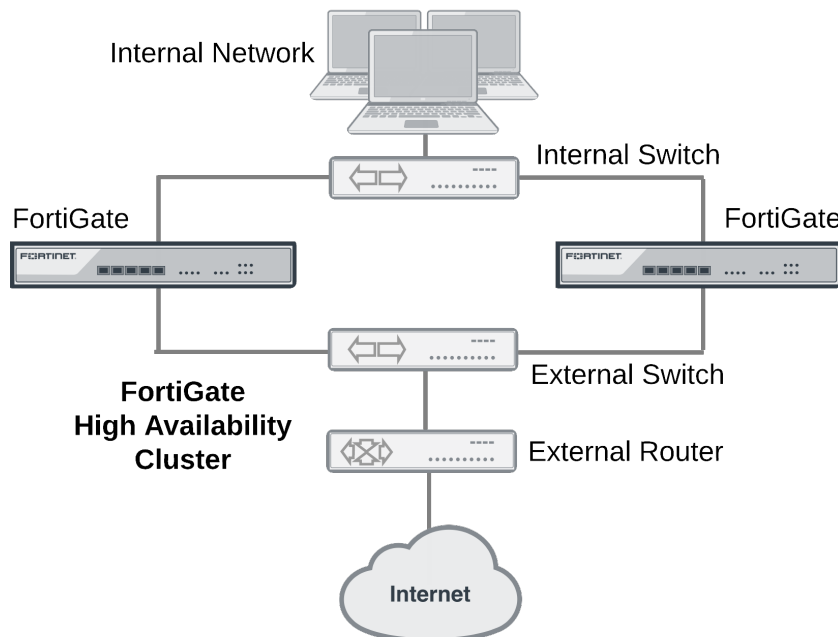
In addition the cluster units must be able to communicate with each other through their heartbeat interfaces. This heartbeat communication is required for the cluster to be created and to continue operating. Without it, the cluster acts like a collection of standalone FortiGates.

On startup, after configuring the cluster units with the same HA configuration and connecting their heartbeat interfaces, the cluster units use the FortiGate Clustering Protocol (FGCP) to find other FortiGates configured for HA operation and to negotiate to create a cluster. During cluster operation, the FGCP shares communication and synchronization information among the cluster units over the heartbeat interface link. This communication and synchronization is called the FGCP heartbeat or the HA heartbeat. Often, this is shortened to just heartbeat.

The cluster uses the FGCP to select the primary unit, and to provide device, link and session failover. The FGCP also manages the two HA modes; active-passive (failover HA) and active-active (load balancing HA).

About the FGCP

FortiGate HA is implemented by configuring two or more FortiGates to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate, processing network traffic and providing normal security services such as firewalling, security profile services, and VPN services.

HA cluster installed between an internal network and the Internet

Inside the cluster the individual FortiGates are called cluster units. These cluster units share state and configuration information. If one cluster unit fails, the other units in the cluster automatically replace that unit, taking over the work that the failed unit was doing. After the failure, the cluster continues to process network traffic and provide normal FortiGate services with virtually no interruption.

Every FortiGate cluster contains one primary unit (also called the master unit) and one or more subordinate units (also called slave or backup units). The primary unit controls how the cluster operates. The role that the subordinate units play depends on the mode in which the cluster operates: (Active-Passive (AP) or Active-Active (AA).

The ability of an HA cluster to continue providing firewall services after a failure is called failover. FGCP failover means that your network does not have to rely on one FortiGate to continue functioning. You can install additional units and form an HA cluster.

A second HA feature, called load balancing, can be used to increase performance. A cluster of FortiGates can increase overall network performance by sharing the load of processing network traffic and providing security services. The cluster appears to your network to be a single device, adding increased performance without changing your network configuration.

Virtual clustering extends HA features to provide failover protection and load balancing for Virtual Domains (VDOMs). See [Virtual clusters on page 143](#).

FortiGate models that support redundant interfaces can be configured to support full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster. For details about full mesh HA, see [Full mesh HA on page 163](#).

FGCP failover protection

The FGCP provides IP/MAC takeover for failover protection by assigning virtual MAC addresses to the primary cluster unit and then sending gratuitous ARP packets from the primary unit interfaces to reprogram the network.

Failover times can be less than a second under optimal conditions. You can fine tune failover performance for your network by adjusting cluster status checking timers, routing table update timers, and wait timers.

An HA cluster fails over if the primary unit fails (a device failure) or experiences a link failure. The cluster can detect link failures for connections to the primary unit using port monitoring and for connections between downstream network components using remote IP monitoring. To compensate for a link failover, the cluster maintains active links to keep traffic flowing between high-priority networks. Port and remote IP monitoring can be fine tuned without disrupting cluster operation.

Session Failover

FGCP session failover maintains TCP, SIP and IPsec VPN sessions after a failure. You can also configure session failover to maintain UDP and ICMP sessions. Session failover does not failover multicast, or SSL VPN sessions. Session failover may not be required for all networks because many TCP/IP, UDP, and ICMP protocols can resume sessions on their own. Supporting session failover adds extra overhead to cluster operations and can be disabled to improve cluster performance if it is not required.

Load Balancing

Active-active HA load balances resource-intensive security profile features such as virus scanning, web filtering, intrusion protection, application control, email filtering and data leak prevention operations among all cluster units to provide better performance than a standalone FortiGate. If network traffic consists of mainly TCP sessions, the FGCP can also load balance all TCP sessions to improve TCP performance in some network configurations. On some FortiGate models you can also load balance UDP sessions. NP4 and NP6 offloading can accelerate HA load balancing (especially TCP and UDP load balancing). HA load balancing schedules can be adjusted to optimize performance for the traffic mix on your network. Weighted load balancing can be used to control the relative amount of sessions processed by each cluster unit.

Virtual Clustering

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGates operating with multiple VDOMS enabled. Not only does virtual clustering provide failover protection for a multiple VDOM configuration, but a virtual cluster can load balance traffic between the cluster units. Load balancing with virtual clustering is quite efficient and load balances all traffic. It is possible to fine tune virtual clustering load balancing in real time to actively optimize load sharing between the cluster units without affecting the smooth operation of the cluster.

Full Mesh HA

High availability improves the reliability of a network by replacing a single point of failure (a single FortiGate) with a cluster that can maintain network traffic if one of the cluster units fails. However, in a normal cluster configuration, single points of failure remain. Full mesh HA removes these single points of failure by allowing you to connect redundant switches to each cluster interface. Full mesh HA is achieved by configuring 802.3ad aggregate or redundant interfaces on the FortiGate and connecting redundant switches to these interfaces. Configuration is a relatively simple extension of the normal aggregate/redundant interface and HA configurations.

Cluster Management

FortiOS HA provides a wide range of cluster management features:

- Automatic continuous configuration synchronization. You can get a cluster up and running almost as quickly as a standalone FortiGate by performing a few basic steps to configure HA settings and minimal network settings on each cluster unit. When the cluster is operating you can configure FortiGate features such as firewalling, content inspection, and VPN in the same way as for a standalone FortiGate. All configuration changes (even complex changes such as switching to multiple VDOM mode or from NAT/Route to Transparent mode) are synchronized among all cluster units.
- Firmware upgrades/downgrades. Upgrading or downgrading cluster firmware is similar to upgrading or downgrading standalone FortiGate firmware. The Firmware is uploaded once to the primary unit and the cluster automatically upgrades or downgrades all cluster units in one operation with minimal or no service interruption.
- Individual cluster unit management. In some cases you may want to manage individual cluster units. You can do so from cluster CLI by navigating to each cluster unit. You can also use the reserved management interface feature to give each cluster unit its own IP address and default route. You can use the reserved management interfaces and IP addresses to connect to the GUI and CLI of each cluster unit and configure an SNMP server to poll each cluster unit.
- Removing and adding cluster units. In one simple step any unit (even the primary unit) can be removed from a cluster and given a new IP address. The cluster keeps operating as it was; the transition happening without interrupting cluster operation. A new unit can also be added to an operating cluster without disrupting network traffic. All you have to do is connect the new unit and change its HA configuration to match the cluster's. The cluster automatically finds and adds the unit and synchronizes its configuration with the cluster.
- Debug and diagnose commands. An extensive range of debug and diagnose commands can be used to report on HA operation and find and fix problems.
- Logging and reporting. All cluster units can be configured to record all log messages. These message can be stored on the individual cluster units or sent to a FortiAnalyzer unit. You can view all cluster unit log messages by logging into any cluster unit.
- FortiManager support. FortiManager understands FortiOS HA and automatically recognizes when you add a FortiOS cluster to the FortiManager configuration.

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. This means that in most cases you only have to make a configuration change once to have it synchronized to all cluster units.

Some configuration settings are not synchronized to support some aspects of FortiGate operation. The following settings are not synchronized among cluster units:

- The FortiGate host name. Allows you to identify cluster units.
- HA override.
- HA device priority.
- Virtual cluster 1 and Virtual cluster 2 device priorities.
- The HA priority (`ha-priority`) setting for a ping server or dead gateway detection configuration.
- The system interface settings of the FortiGate interface that becomes the HA reserved management interface.
- The default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.
- The dynamic weighted load balancing thresholds and high and low watermarks.

Synchronizing the configuration (and settings that are not synchronized)

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. This means that in most cases you only have to make a configuration change once to have it synchronized to all cluster units. This includes special configuration settings that include extra information (for example, 3rd party certificates, replacement message text files and graphics and so on).

Some configuration settings are not synchronized to support some aspects of FortiGate operation. The following settings are not synchronized among cluster units:

- The FortiGate host name. Allows you to identify cluster units.
- HA override.
- HA device priority.
- Virtual cluster 1 and Virtual cluster 2 device priorities.
- The HA priority (`ha-priority`) setting for a ping server or dead gateway detection configuration.
- The system interface settings of the FortiGate interface that becomes the HA reserved management interface.
- The default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.
- The dynamic weighted load balancing thresholds and high and low watermarks.

In addition licenses are not synchronized since each FortiGate must be licensed separately. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMS).

Preparing the FortiGates before setting up a FGCP cluster

Before creating an FGCP cluster you should complete the following setup on each FortiGate.

DHCP and PPPoE

Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP or PPPoE you should temporarily switch it to a static address and enable DHCP or PPPoE after the cluster has been established.

Firmware version

Make sure the FortiGates are running the same FortiOS firmware version.

About HA and licensing

All of the FortiGates in a cluster must have the same level of licensing. This includes FortiGuard, FortiCloud, FortiClient, VDOMs (if applicable) and FortiOS Carrier (if applicable).

If one of the FortiGates in a cluster has a lower level of licensing than other FortiGates in the cluster, then all of the FortiGates in the cluster will revert to that lower licensing level. For example, if you only purchase FortiGuard

Web Filtering for one of the FortiGates in a cluster, when the cluster is operating, none of the cluster units will support FortiGuard Web Filtering.

An exception is FortiToken licensing. FortiToken activations are completed one FortiGate unit and synchronized to all of the FortiGates in the cluster.

FortiOS Carrier license

If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license. All FortiGates in the cluster must be licensed for FortiOS Carrier.

Support contracts and FortiGuard, FortiCloud, FortiClient, VDOMs Licensing

Register and apply these licenses to each FortiGate. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMS). All FortiGates in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient and VDOMs.

License Information

	Support Contract	Registration	✓ Registered (bdickie@fortinet.com)	Launch Portal
	FortiGuard	IPS & Application Control	✓ Licensed (Expires 2016-08-22)	
		AntiVirus	✓ Licensed (Expires 2016-08-22)	
		Web Filtering	✓ Licensed (Expires 2016-08-21)	
		Anti-Spam Filtering	✓ Licensed (Expires 2016-08-21)	
	FortiCloud	Account		Activate
	FortiSandbox	FortiSandbox Appliance	✗ Not Configured	Configure
	FortiClient	Status	✓ Free License	How to Purchase
		Clients Registered	0 of 10	Enter License
		FortiClient Installers		Details
	FortiToken Mobile	Tokens Assigned	0 of 2	

FortiToken Licenses

You only need one set of FortiToken licenses for the HA cluster and you only need to activate each token once. Normally you would activate your tokens on the primary unit and this configuration and the seed information will be synchronized to all cluster members so all tokens will then be activated for all cluster members.

If you have added FortiToken licenses and activated FortiTokens on a standalone FortiGate unit before configuring HA the licenses and the FortiToken activations will usually be synchronized to all cluster units after forming a cluster. To make sure this goes smoothly you can make sure the FortiGate that you have added the

licenses to become the primary unit when setting up the cluster as described in [How to set up FGCP clustering \(recommended steps\)](#) on page 66.

Certificates

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

Configuring FortiGates for FGCP HA operation

Each FortiGate in the cluster must have the same HA configuration. Once the cluster is connected, you can configure it in the same way as you would configure a standalone FortiGate. The following example sets the HA mode to active-passive and the HA password to HA_pass.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP or PPPoE you should temporarily switch it to a static address and enable DHCP or PPPoE after the cluster has been established.

Make sure both FortiGates are running the same FortiOS firmware version. Register and apply licenses to both FortiGates before adding them to the cluster. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information

Support Contract
Registration

✓ Registered (bdickie@fortinet.com)

Launch Portal

FortiGuard

IPS & Application Control

✓ Licensed (Expires 2016-08-22)

AntiVirus

✓ Licensed (Expires 2016-08-22)

Web Filtering

✓ Licensed (Expires 2016-08-21)

Anti-Spam Filtering

✓ Licensed (Expires 2016-08-21)

FortiCloud

Account

Activate

FortiSandbox

FortiSandbox Appliance

✗ Not Configured

Configure

FortiClient

Status

✓ Free License

How to Purchase

Enter License

Clients Registered

0 of 10

Details

FortiClient Installers

FortiToken Mobile

Tokens Assigned

0 of 2

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

32

High Availability for FortiOS 5.4.4
 Fortinet Technologies Inc.

We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.

To configure a FortiGate for HA operation - GUI

1. Power on the FortiGate to be configured.
2. Log into the GUI.
3. On the Dashboard **System Information** dashboard widget, beside **Host Name** select **Change**.
4. Enter a new Host Name for this FortiGate.
Changing the host name makes it easier to identify individual cluster units when the cluster is operating.
5. Go to **System > HA** and change the following settings:

Mode	Active-Passive
Group Name	Example_cluster
Password	HA_pass
The password must be the same for all FortiGates in the cluster.	

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

6. Select **OK**.
The FortiGate negotiates to establish an HA cluster. When you select **OK** you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all ARP table entries). You may be able to delete the ARP table of your management PC from a command prompt using a command similar to `arp -d`.
7. Power off the FortiGate.
8. Repeat this procedure for all of the FortiGates in the cluster.
Once all of the units are configured, continue by connecting the FortiGate HA cluster below.

To configure a FortiGate for HA operation - CLI

1. Power on the FortiGate to be configured.
2. Log into the CLI.
3. Enter the following command to change the FortiGate host name.

```
config system global
  set hostname Example1_host
end
```

Changing the host name makes it easier to identify individual cluster units when the cluster is operating.

4. Enter the following command to enable HA:

```
config system ha
  set mode active-passive
  set group-name Example_cluster
  set password HA_pass
end
```

You can accept the default configuration for the remaining HA options and change them later, once

the cluster is operating.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces. To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

5. Power off the FortiGate.
6. Repeat this procedure for all of the FortiGates in the cluster.
Once all of the units are configured, continue with connecting the FortiGate HA cluster.

Connecting a FortiGate HA cluster

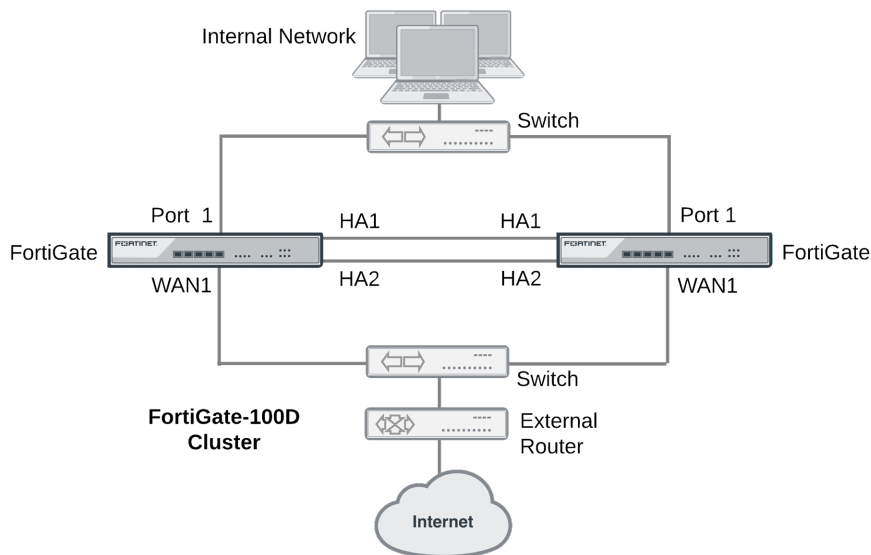
Use the following procedure to connect a cluster. Connect the cluster units to each other and to your network. You must connect all matching interfaces in the cluster to the same switch, then connect these interfaces to their networks using the same switch.

Although you can use hubs, Fortinet recommends using switches for all cluster connections for the best performance.

Connecting an HA cluster to your network temporarily interrupts communications on the network because new physical connections are being made to route traffic through the cluster. Also, starting the cluster interrupts network traffic until the individual cluster units are functioning and the cluster completes negotiation. Cluster negotiation is automatic and normally takes just a few seconds. During system startup and negotiation all network traffic is dropped.

This section describes how to connect the cluster shown below, which consists of two FortiGate-100D units to be connected between the Internet and a head office internal network. The wan1 interfaces of the FortiGate connect the cluster to the Internet and the internal interfaces connect the cluster to the internal network. The ha1 and ha2 interfaces are used for redundant HA heartbeat links.

Example cluster connections



To connect a FortiGate HA cluster

1. Connect the WAN1 interfaces of each cluster unit to a switch connected to the Internet.
2. Connect the Port1 interfaces of each cluster unit to a switch connected to the internal network.
3. Connect the HA1 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)
4. Connect the HA2 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)
5. Power on both of the FortiGates.

As the cluster units start, they negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally just takes a few seconds.

At least one heartbeat interface should be connected together for the cluster to operate.

Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.


You could use one switch to connect all four heartbeat interfaces. However, this is not recommended because if the switch fails both heartbeat interfaces will become disconnected.

6. You can now configure the cluster as if it is a single FortiGate.

Verifying the cluster status from the Unit Operation dashboard widget

The Unit Operation dashboard widget includes the serial number and hostname of all of the FortiGates in the cluster as well as an indication of the sync status of each cluster member.

Unit Operation



HA Cluster Members	Sync Status	Role
FG-5KD3914800344/FG-5KD3914800344	✓	Master
FG-5KD3914800284/FG-5KD3914800284	✓	Slave

Active-passive and active-active HA

The first decision to make when configuring FortiGate HA is whether to choose active-passive or active-active HA mode. To configure the HA mode, go to **System > HA** and set Mode to **Active-Passive** or **Active-Active**.

From the CLI enter the following command to set the HA mode to active-passive:

```
config system ha
  set mode a-p
end
```

To form a cluster, all cluster units must be set to the same mode. You can also change the mode after the cluster is up and running. Changing the mode of a functioning cluster causes a slight delay while the cluster renegotiates to operate in the new mode and possibly select a new primary unit.

Active-passive HA (failover protection)

An active-passive (A-P) HA cluster provides hot standby failover protection. An active-passive cluster consists of a primary unit that processes communication sessions, and one or more subordinate units. The subordinate units are connected to the network and to the primary unit but do not process communication sessions. Instead, the subordinate units run in a standby state. In this standby state, the configuration of the subordinate units is synchronized with the configuration of the primary unit and the subordinate units monitor the status of the primary unit.

Active-passive HA provides transparent device failover among cluster units. If a cluster unit fails, another immediately take its place.

Active-passive HA also provides transparent link failover among cluster units. If a cluster unit interface fails or is disconnected, this cluster unit updates the link state database and the cluster negotiates and may select a new primary unit.

If session failover (also called session pickup) is enabled, active-passive HA provides session failover for some communication sessions.

The following example shows how to configure a FortiGate for active-passive HA operation. You would enter the exact same commands on every FortiGate in the cluster.

```
config system ha
  set mode a-p
  set group-name myname
  set password HApass
end
```

Active-active HA (load balancing and failover protection)

By default, active-active HA load balancing distributes proxy-based security profile processing to all cluster units. Proxy-based security profile processing is CPU and memory-intensive, so FGCP load balancing may result in higher throughput because resource-intensive processing is distributed among all cluster units.

Normally, sessions accepted by policies that don't include security profiles are not load balanced and are processed by the primary unit. You can configure active-active HA to load balance additional sessions.

An active-active HA cluster consists of a primary unit that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. In an active-active cluster the subordinate units are also considered active since they also process content processing sessions. In all other ways active-active HA operates the same as active-passive HA.

The following example shows how to configure a FortiGate for active-active HA operation. You would enter the exact same commands on every FortiGate in the cluster.

```
config system ha
  set mode a-a
  set group-name myname
  set password HApass
end
```

Identifying the cluster and cluster units

You can use the cluster group name, group id, and password to identify a cluster and distinguish one cluster from another. If you have more than one cluster on the same network, each cluster must have a different group name, group id, and password.

Group name

Use the group name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating,

you can change the group name. The group name change is synchronized to all cluster units. The group name appears on the FortiGate dashboard of a functioning cluster as the **Cluster Name**.

To add or change the group name from the GUI go to **System > HA** and change the **Group Name**.

Enter the following CLI command to change the group name to Cluster_name:

```
config system ha
    set group-name Cluster_name
end
```

Password

Use the password to identify the cluster. You should always change the password when configuring a cluster. The password must be the same for all FortiGates before they can form a cluster. When the cluster is operating you can change the password, if required. Two clusters on the same network cannot have the same password.

To change the password from the GUI go to **System > HA** and change the **Password**.

Enter the following CLI command to change the password to ha_pwd:

```
config system ha
    set password ha_pwd
end
```

Group ID

Similar to the group name, the group ID also identifies the cluster. In most cases you do not have to change the group ID. However, you should change the group ID if you have more than one cluster on the same network. All members of the HA cluster must have the same group ID. The group ID is a number from 0 to 255.

Changing the group ID changes the cluster virtual MAC address. If two clusters on the same network have the same group ID you may encounter MAC address conflicts.

Enter the following CLI command to change the group ID to 10:

```
config system ha
    set group-id 10
end
```

Device failover, link failover, and session failover

The FGCP provides transparent device and link failover. You can also enable session pickup to provide session failover. A failover can be caused by a hardware failure, a software failure, or something as simple as a network cable being disconnected causing a link failover. When a failover occurs, the cluster detects and recognizes the failure and takes steps to respond so that the network can continue to operate without interruption. The internal operation of the cluster changes, but network components outside of the cluster notice little or no change.

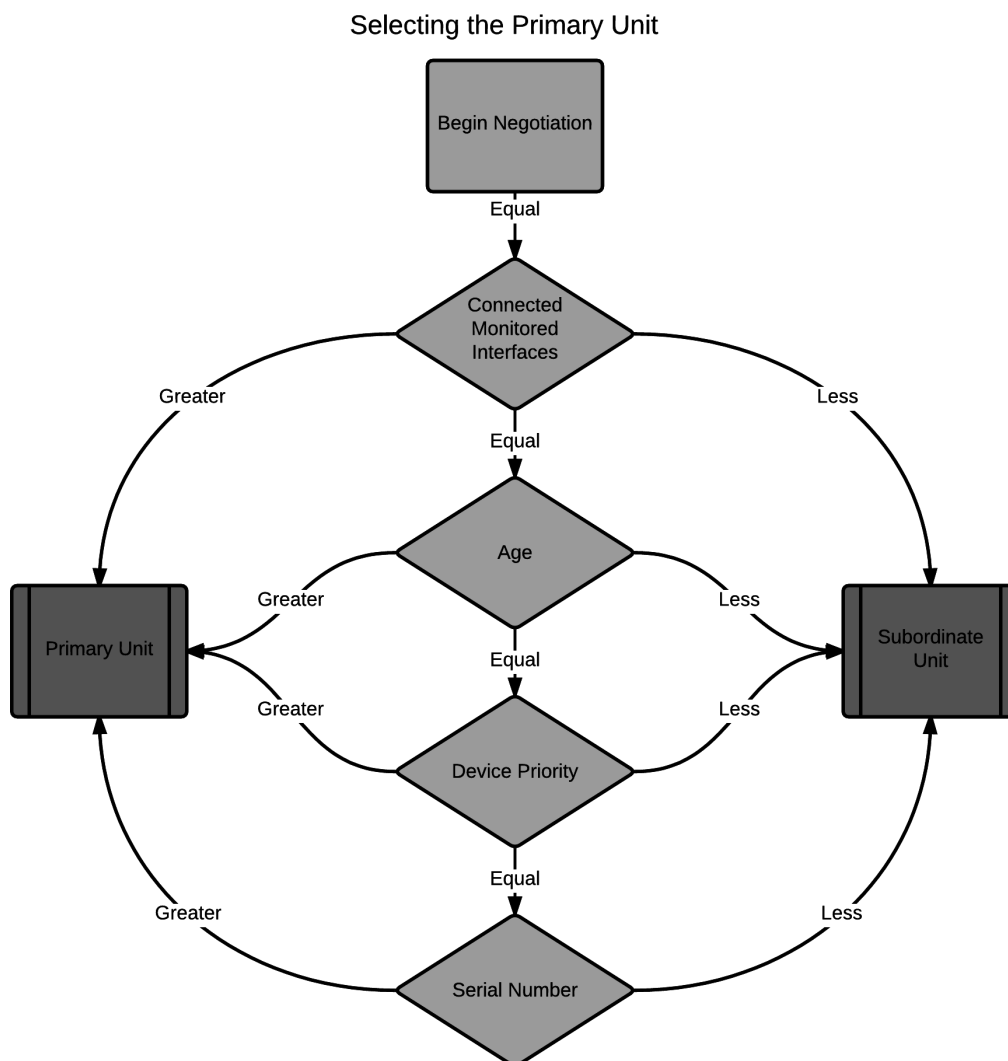
If a failover occurs, the cluster also records log messages about the event and can be configured to send log messages to a syslog server and to a FortiAnalyzer unit. The cluster can also send SNMP traps and alert email messages. These alerts can notify network administrators of the failover and may contain information that the network administrators can use to find and fix the problem that caused the failure.

For a complete description of device failover, link failover, and session failover, how clusters support these types of failover, and how FortiGate HA clusters compensate for a failure to maintain network traffic flow see [HA and failover protection on page 215](#).

Primary unit selection

Once FortiGates recognize that they can form a cluster, the cluster units negotiate to select a primary unit. Primary unit selection occurs automatically based on the criteria shown below. After the cluster selects the primary unit, all of the remaining cluster units become subordinate units.

Negotiation and primary unit selection also takes place if a primary unit fails (device failover) or if a monitored interface fails or is disconnected (link failover). During a device or link failover, the cluster renegotiates to select a new primary unit also using the criteria shown below.



For many basic HA configurations primary unit selection simply selects the cluster unit with the highest serial number to become the primary unit. A basic HA configuration involves setting the HA mode to active-passive or active-active and configuring the cluster group name and password. Using this configuration, the cluster unit with the highest serial number becomes the primary unit because primary unit selection disregards connected monitored interfaces (because interface monitoring is not configured), the age of the cluster units would usually always be the same, and all units would have the same device priority.

Using the serial number is a convenient way to differentiate cluster units; so basing primary unit selection on the serial number is predictable and easy to understand and interpret. Also the cluster unit with the highest serial number would usually be the newest FortiGate with the most recent hardware version. In many cases you may not need active control over primary unit selection, so basic primary unit selection based on serial number is sufficient.

In some situations you may want have control over which cluster unit becomes the primary unit. You can control primary unit selection by setting the device priority of one cluster unit to be higher than the device priority of all other cluster units. If you change one or more device priorities, during negotiation, the cluster unit with the highest device priority becomes the primary unit. As shown above, the FGCP selects the primary unit based on device priority before serial number. For more information about how to use device priorities, see [Primary unit selection and device priority on page 44](#).

The only other way that you can influence primary unit selection is by configuring interface monitoring (also called port monitoring). Using interface monitoring you can make sure that cluster units with failed or disconnected monitored interfaces cannot become the primary unit. See [Primary unit selection and monitored interfaces on page 41](#).

Finally, the age of a cluster unit is determined by a number of operating factors. Normally the age of all cluster units is the same so normally age has no effect on primary unit selection. Age does affect primary unit selection after a monitored interface failure. For more information about age, see [Primary unit selection and age on page 41](#).

Viewing how the primary unit was selected

You can use the `get system ha status` command to see how the primary unit was selected. The output of this command contains a section called `Master selected using` that shows a history of how the primary unit was selected. For example, when a cluster first forms this part of the command output could have one line showing that the primary unit is the cluster unit with the highest uptime.

```
get system ha status
.
.
.
Master selected using:
    <2016/10/12 11:13:23> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
.
.
.
```

Over time more messages could be added as the cluster negotiates to choose a new primary unit on different occasions. The command output below shows the cluster negotiated four times over a few days.

```
get system ha status
.
.
.
```



```
Master selected using:
  <2016/10/16 11:36:07> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
  <2016/10/15 11:24:11> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
  <2016/10/13 11:15:13> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
  <2016/10/11 11:13:23> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
.
.
.
```

Primary unit selection and monitored interfaces

If you have configured interface monitoring the cluster unit with the highest number of monitored interfaces that are connected to networks becomes the primary unit. Put another way, the cluster unit with the highest number of failed or disconnected monitored interfaces cannot become the primary unit.

Normally, when a cluster starts up, all monitored interfaces of all cluster units are connected and functioning normally. So monitored interfaces do not usually affect primary unit selection when the cluster first starts.

A cluster always renegotiates when a monitored interface fails or is disconnected (called link failover). A cluster also always renegotiates when a failed or disconnected monitored interface is restored.

If a primary unit monitored interface fails or is disconnected, the cluster renegotiates and if this is the only failed or disconnected monitored interface the cluster selects a new primary unit.

If a subordinate unit monitored interface fails or is disconnected, the cluster also renegotiates but will not necessarily select a new primary unit. However, the subordinate unit with the failed or disconnected monitored interface cannot become the primary unit.

Multiple monitored interfaces can fail or become disconnected on more than one cluster unit. Each time a monitored interface is disconnected or fails, the cluster negotiates to select the cluster unit with the most connected and operating monitored interfaces to become the primary unit. In fact, the intent of the link failover feature is just this, to make sure that the primary unit is always the cluster unit with the most connected and operating monitored interfaces.

Primary unit selection and age

The cluster unit with the highest age value becomes the primary unit. The age of a cluster unit is the amount of time since a monitored interface failed or is disconnected. Age is also reset when a cluster unit starts (boots up). So, when all cluster units start up at about the same time, they all have the same age. Age does not affect primary unit selection when all cluster units start up at the same time. Age also takes precedence over priority for primary unit selection.

If a link failure of a monitored interface occurs, the age value for the cluster unit that experiences the link failure is reset. So, the cluster unit that experienced the link failure also has a lower age value than the other cluster units. This reduced age does not effect primary unit selection because the number of link failures takes precedence over the age.

If the failed monitored interface is restored the cluster unit that had the failed monitored interface cannot become the primary unit because its age is still lower than the age of the other cluster units.

In most cases, the way that age is handled by the cluster reduces the number of times the cluster selects a new primary unit, which results in a more stable cluster since selecting a new primary unit has the potential to disrupt traffic.

Cluster age difference margin (grace period)

In any cluster, some of the cluster units may take longer to start up than others. This startup time difference can happen as a result of a number of issues and does not affect the normal operation of the cluster. To make sure that cluster units that start slower can still become primary units, by default the FGCP ignores age differences of up to 5 minutes (300 seconds).

In most cases, during normal operation this age difference margin or grace period helps clusters function as expected. However, the age difference margin can result in some unexpected behavior in some cases:

- During a cluster firmware upgrade with `uninterruptible-upgrade` enabled (the default configuration) the cluster should not select a new primary unit after the firmware of all cluster units has been updated. But since the age difference of the cluster units is most likely less than 300 seconds, age is not used to affect primary unit selection and the cluster may select a new primary unit.
- During failover testing where cluster units are failed over repeatedly the age difference between the cluster units will most likely be less than 5 minutes. During normal operation, if a failover occurs, when the failed unit rejoins the cluster its age will be very different from the age of the still operating cluster units so the cluster will not select a new primary unit. However, if a unit fails and is restored in a very short time the age difference may be less than 5 minutes. As a result the cluster may select a new primary unit during some failover testing scenarios.

Changing the cluster age difference margin

You can change the cluster age difference margin using the following command:

```
config system ha
    set ha-uptime-diff-margin 60
end
```

This command sets the cluster age difference margin to 60 seconds (1 minute). The age difference margin range is 1 to 65535 seconds. The default is 300 seconds.

You may want to reduce the margin if during failover testing you don't want to wait the default age difference margin of 5 minutes. You may also want to reduce the margin to allow uninterruptible upgrades to work. See [Operating clusters and virtual clusters on page 177](#).

You may want to increase the age margin if cluster unit startup time differences are larger than 5 minutes.

Displaying cluster unit age differences

You can use the CLI command `diagnose sys ha dump-by all-vcluster` to display the age difference of the units in a cluster. This command also displays information about a number of HA-related parameters for each cluster unit. You can enter the command from the primary unit CLI or you can enter the command from a subordinate unit after using `execute ha manage` to log into a subordinate unit CLI. The information displayed by the command is relative to the unit that you enter the command from.

For example, a cluster of two FortiGate-5001C units with no changes to the default HA configuration except to enable port monitoring for port1. Entering the `diagnose sys ha dump-by all-vcluster` command from the primary unit CLI displays information similar to the following:

```
diagnose sys ha dump-by all-vcluster
    HA information.
vcluster id=1, nentry=2, state=work, digest=4.e8.62.17.7b.1d...
ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,
```

```

    flag=0x01,time=0,mon=0
    mondev=port1,50
    ventry idx=1,id=1,FG-5KC3E13800051,prio=128,0,claimed=0,override=0,
    flag=0x00,time=189,mon=0

```

The command displays one `ventry` line for each cluster unit. The first `ventry` in the example contains information for the cluster unit that you are logged into (usually the primary unit). The other `ventry` lines contain information for the other units in the cluster (in the example there is only one other cluster unit). The command also includes a `mondev` entry that displays the interface monitoring configuration.

The `time` field is always 0 for the unit that you are logged into. The `time` field for the other cluster unit is the age difference between the unit that you are logged into and the other cluster unit. The age difference is in the form seconds/10.

In the example, the age of the subordinate unit is 18.9 seconds more than the age of the primary unit. The age difference is less than 5 minutes (less than 300 seconds) so age has no effect on primary unit selection. The cluster selected the unit with the highest serial number to be the primary unit.

If you use `execute ha manage 1` to log into the subordinate unit CLI and enter `diagnose sys ha dump 1` you get results similar to the following:

```

diagnose sys ha dump-by all-vcluster
    HA information.
    vcluster id=1, nentry=2, state=standby, digest=4.e8.62.17.7b.1d...
    ventry idx=1,id=1,FG-5KC3E13800051,prio=128,0,claimed=0,override=0,
    flag=0x01,time=0,mon=0
    mondev=port1,50
    ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=1,override=0,
    flag=0x00,time=-189,mon=0

```

The `time` for the primary unit is -189, indicating that age of the subordinate unit age is 18.9 seconds higher than the primary unit age.

If port1 (the monitored interface) of the primary unit is disconnected, the cluster renegotiates and the former subordinate unit becomes the primary unit. When you log into the new primary unit CLI and enter `diagnose sys ha dump-by all-vcluster` you could get results similar to the following:

```

diagnose sys ha dump-by all-vcluster
    HA information.
    vcluster id=1, nentry=2, state=work, digest=3.f8.d1.63.4d.d2...
    ventry idx=0,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,
    override=0,flag=1,time=0,mon=0
    mondev=port1,50
    ventry idx=1,id=1,FG-5KC3E13800084,prio=128,-50,claimed=0,
    override=0,flag=0,time=1362,mon=0

```

The command results show that the age of the new primary unit is 136.2 seconds higher than the age of the new subordinate unit.

If port1 of the former primary unit is reconnected the cluster will once again make this the primary unit because the age difference will still be less than 300 seconds. When you log into the primary unit CLI and enter `diagnose sys ha dump-by all-vcluster` you get results similar to the following:

```

diagnose sys ha dump-by all-vcluster
    HA information.
    vcluster id=1, nentry=2, state=work, digest=4.a5.60.11.cf.d4...
    ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,
    override=0,flag=1,time=0,mon=0
    mondev=port1,50
    ventry idx=1,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,
    override=0,flag=0,time=-1362,mon=0

```

Resetting the age of all cluster units

In some cases, age differences among cluster units can result in the wrong cluster unit or the wrong virtual cluster becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units when it rejoins the cluster. Since age takes precedence over priority, the priority of this cluster unit will not be a factor in primary unit selection.

This problem also affects virtual cluster VDOM partitioning in a similar way. After a reboot of one of the units in a virtual cluster configuration, traffic for all VDOMs could continue to be processed by the cluster unit that did not reboot. This can happen because the age of both virtual clusters on the unit that did not reboot is greater than the age of both virtual clusters on the unit that rebooted.

One way to resolve this issue is to reboot all of the cluster units at the same time so that the age of all of the cluster units is reset. However, rebooting cluster units may interrupt or at least slow down traffic. If you would rather not reboot all of the cluster units you can instead use the following command to reset the age of individual cluster units.

```
diagnose sys ha reset-uptime
```

This command resets the age of a unit back to zero so that if no other unit in the cluster was reset at the same time, it will now have the lowest age. You would use this command to reset the age of the cluster unit that is currently the primary unit. Since it will have the lowest age, the other unit in the cluster will have the highest age and can then become the primary unit.



The `diagnose sys ha reset-uptime` command should only be used as a temporary solution. The command resets the HA age internally and does not affect the up time displayed for cluster units using the `diagnose sys ha dump-by all-vcluster` command or the up time displayed on the Dashboard or cluster members list. To make sure the actual up time for cluster units is the same as the HA age you should reboot the cluster units during a maintenance window.

Primary unit selection and device priority

A cluster unit with the highest device priority becomes the primary unit when the cluster starts up or renegotiates. By default, the device priority for all cluster units is 128. You can change the device priority to control which FortiGate becomes the primary unit during cluster negotiation. All other factors that influence primary unit selection either cannot be configured (age and serial number) or are synchronized among all cluster units (interface monitoring). You can set a different device priority for each cluster unit. During negotiation, if all monitored interfaces are connected, and all cluster units enter the cluster at the same time (or have the same age), the cluster with the highest device priority becomes the primary unit.

A higher device priority does not affect primary unit selection for a cluster unit with the most failed monitored interfaces or with an age that is higher than all other cluster units because failed monitored interfaces and age are used to select a primary unit before device priority.

Increasing the device priority of a cluster unit does not always guarantee that this cluster unit will become the primary unit. During cluster operation, an event that may affect primary unit selection may not always result in the cluster renegotiating. For example, when a unit joins a functioning cluster, the cluster will not renegotiate. So if a unit with a higher device priority joins a cluster the new unit becomes a subordinate unit until the cluster renegotiates.



Enabling the `override` HA CLI keyword makes changes in device priority more effective by causing the cluster to negotiate more often to make sure that the primary unit is always the unit with the highest device priority. For more information about `override`, see [Primary unit selection on page 39](#).

Controlling primary unit selection by changing the device priority

You set a different device priority for each cluster unit to control the order in which cluster units become the primary unit when the primary unit fails.

To change the device priority from the GUI go to **System > HA** and change the **Device Priority**.

Enter the following CLI command to change the device priority to 200:

```
config system ha
    set priority 200
end
```

The device priority is not synchronized among cluster units. In a functioning cluster you can change the device priority of any unit in the cluster. Whenever you change the device priority of a cluster unit, when the cluster negotiates, the unit with the highest device priority becomes the primary unit.

The following example shows how to change the device priority of a subordinate unit to 255 so that this subordinate unit becomes the primary unit. You can change the device priority of a subordinate unit by going to **System > HA** and selecting the Edit icon for the subordinate unit. Or from the CLI you can use the `execute ha manage 0` command to connect to the highest priority subordinate unit. After you enter the following commands the cluster renegotiates and selects a new primary unit.

```
execute ha manage 1
    config system ha
        set priority 255
    end
```

If you have three units in a cluster you can set the device priorities as shown below. When the cluster starts up, cluster unit A becomes the primary unit because it has the highest device priority. If unit A fails, unit B becomes the primary unit because unit B has a higher device priority than unit C.

Example device priorities for a cluster of three FortiGates

Cluster unit	Device priority
A	200
B	100
C	50

When configuring HA you do not have to change the device priority of any of the cluster units. If all cluster units have the same device priority, when the cluster first starts up the FGCP negotiates to select the cluster unit with the highest serial number to be the primary unit. Clusters also function normally if all units have the same device priority.

You can change the device priority if you want to control the roles that individual units play in the cluster. For example, if you want the same unit to always become the primary unit, set this unit device priority higher than the

device priority of other cluster units. Also, if you want a cluster unit to always become a subordinate unit, set this cluster unit device priority lower than the device priority of other cluster units.

If you have a cluster of three units you can set a different priority for each unit to control which unit becomes the primary unit when all three cluster units are functioning and which will be the primary unit when two cluster units are functioning.

The device priority range is 0 to 255. The default device priority is 128.

If you are configuring a virtual cluster, if you have added virtual domains to both virtual clusters, you can set the device priority that the cluster unit has in virtual cluster 1 and virtual cluster 2. If a FortiGate has different device priorities in virtual cluster 1 and virtual cluster 2, the FortiGate may be the primary unit in one virtual cluster and the subordinate unit in the other.

Primary unit selection and the FortiGate serial number

The cluster unit with the highest serial number is more likely to become the primary unit. When first configuring FortiGates to be added to a cluster, if you do not change the device priority of any cluster unit, then the cluster unit with the highest serial number always becomes the primary unit.

Age does take precedence over serial number, so if a cluster unit takes longer to join a cluster for some reason (for example if one cluster unit is powered on after the others), that cluster unit will not become the primary unit because the other units have been in the cluster longer.

Device priority and failed monitored interfaces also take precedence over serial number. A higher device priority means a higher priority. So if you set the device priority of one unit higher or if a monitored interface fails, the cluster will not use the FortiGate serial number to select the primary unit.

Points to remember about primary unit selection

Some points to remember about primary unit selection:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored interfaces > Age > Device Priority > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered if a cluster unit fails or if a monitored interface fails.
- If the HA age difference is more than 5 minutes (300 seconds), the cluster unit that is operating longer becomes the primary unit.
- If HA age difference is less than 5 minutes (300 seconds), the device priority and FortiGate serial number selects the cluster unit to become the primary unit.
- Every time a monitored interface fails the HA age of the cluster unit is reset to 0.
- Every time a cluster unit restarts the HA age of the cluster unit is reset to 0.

Temporarily setting a cluster unit to be the primary unit

You can use the following diagnose command to set a cluster unit to be the primary unit.

```
diagnose sys ha set-as-master enable
```



This command is intended for demonstration purposes and not for production use.
This command may not be visible for all FortiOS versions.

When you enter this command, the cluster immediately re-negotiates and the cluster unit on which you entered this command becomes the primary unit. This change is temporary and will be reverted if the cluster unit restarts.

You can also use the following command from the same cluster unit to turn this option off, causing the cluster to renegotiate and select a new primary unit.

```
diagnose sys ha set-as-master disable
```

You can also configure when to disabling the set-as-master setting. For example, to disable the set as master setting on January 25, 2015 you can enter a date after the disable keyword:

```
diagnose sys ha set-as-master disable 2015 01 25
```

HA override

The HA `override` CLI keyword is disabled by default. When `override` is disabled a cluster may not always renegotiate when an event occurs that affects primary unit selection. For example, when `override` is disabled a cluster will not renegotiate when you change a cluster unit device priority or when you add a new cluster unit to a cluster. This is true even if the unit added to the cluster has a higher device priority than any other unit in the cluster. Also, when `override` is disabled a cluster does not negotiate if the new unit added to the cluster has a failed or disconnected monitored interface.



For a virtual cluster configuration, `override` is enabled by default for both virtual clusters when you enable virtual cluster 2. For more information, see [Virtual clusters on page 143](#).

In most cases you should keep `override` disabled to reduce how often the cluster negotiates. Frequent negotiations may cause frequent traffic interruptions.

However, if you want to make sure that the same cluster unit always operates as the primary unit and if you are less concerned about frequent cluster negotiation you can set its device priority higher than other cluster units and enable `override`.

To enable `override`, connect to each cluster unit CLI (using the `execute ha manage` command) and use the `config system ha` CLI command to enable `override`.

For `override` to be effective, you must also set the device priority highest on the cluster unit that you want to always be the primary unit. To increase the device priority, from the CLI use the `config system ha` command and increase the value of the `priority` keyword to a number higher than the default priority of 128.

You can also increase the device priority from the GUI by going to **System > HA**. To increase the device priority of the primary unit select edit for the primary or subordinate unit and set the **Device Priority** to a number higher than 128.



The `override` setting and device priority value are not synchronized to all cluster units. You must enable `override` and adjust device priority manually and separately for each cluster unit.

With `override` enabled, the primary unit with the highest device priority will always become the primary unit. Whenever an event occurs that may affect primary unit selection, the cluster negotiates. For example, when


`override` is enabled a cluster renegotiates when you change the device priority of any cluster unit or when you add a new unit to a cluster.

Override and primary unit selection

Enabling `override` changes the order of primary unit selection. As shown below, if `override` is enabled, primary unit selection considers device priority before age and serial number. This means that if you set the device priority higher on one cluster unit, with `override` enabled this cluster unit becomes the primary unit even if its age and serial number are lower than other cluster units.

Similar to when `override` is disabled, when `override` is enabled primary unit selection checks for connected monitored interfaces first. So if interface monitoring is enabled, the cluster unit with the most disconnected monitored interfaces cannot become the primary unit, even if the unit has the highest device priority.

If all monitored interfaces are connected (or interface monitoring is not enabled) and the device priority of all cluster units is the same then age and serial number affect primary unit selection.



Controlling primary unit selection using device priority and override

To configure one cluster unit to always become the primary unit you should set its device priority to be higher than the device priorities of the other cluster units and you should enable `override` on all cluster units.

Using this configuration, when the cluster is operating normally the primary unit is always the unit with the highest device priority. If the primary unit fails the cluster renegotiates to select another cluster unit to be the primary unit. If the failed primary unit recovers, starts up again and rejoins the cluster, because `override` is enabled, the cluster renegotiates. Because the restarted primary unit has the highest device priority it once again becomes the primary unit.

In the same situation with `override` disabled, because the age of the failed primary unit is lower than the age of the other cluster units, when the failed primary unit rejoins the cluster it does not become the primary unit. Instead, even though the failed primary unit may have the highest device priority it becomes a subordinate unit because its age is lower than the age of all the other cluster units.

Points to remember about primary unit selection when override is enabled

Some points to remember about primary unit selection when `override` is enabled:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored Interfaces > Device Priority > Age > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered whenever an event occurs which may affect primary unit selection. For example negotiation occurs, when you change the device priority, when you add a new unit to a cluster, if a cluster unit fails, or if a monitored interface fails.
- Device priority is considered before age. Otherwise age is handled the same when `override` is enabled.

Configuration changes can be lost if override is enabled

In some cases, when `override` is enabled and you make configuration changes to an HA cluster these changes can be lost. For example, consider the following sequence:

1. A cluster of two FortiGates is operating with `override` enabled.
 - FGT-A: Primary unit with device priority 200 and with `override` enabled
 - FGT-B: Subordinate unit with device priority 100 and with `override` disabled
 - If both units are operating, FGT-A always becomes the primary unit because FGT-A has the highest device priority.
2. FGT-A fails and FGT-B becomes the new primary unit.
3. The administrator makes configuration changes to the cluster.

The configuration changes are made to FGT-B because FGT-B is operating as the primary unit. These configuration changes are not synchronized to FGT-A because FGT-A is not operating.
4. FGT-A is restored and starts up again.
5. The cluster renegotiates and FGT-A becomes the new primary unit.
6. The cluster recognizes that the configurations of FGT-A and FGT-B are not the same.
7. The configuration of FGT-A is synchronized to FGT-B.

The configuration is always synchronized from the primary unit to the subordinate units.
8. The cluster is now operating with the same configuration as FGT-A. The configuration changes made to FGT-B have been lost.

The solution

When `override` is enabled, you can prevent configuration changes from being lost by doing the following:

- Verify that all cluster units are operating before making configuration changes (from the GUI go to **System > HA** to view the cluster members list or from the FortiOS CLI enter `get system ha status`).
- Make sure the device priority of the primary unit is set higher than the device priorities of all other cluster units before making configuration changes.
- Disable `override` either permanently or until all configuration changes have been made and synchronized to all cluster units.

Override and disconnecting a unit from a cluster

A similar scenario to that described above may occur when `override` is enabled and you use the Disconnect from Cluster option from the GUI or the `execute ha disconnect` command from the CLI to disconnect a cluster unit from a cluster.

Configuration changes made to the cluster can be lost when you reconnect the disconnected unit to the cluster. You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. Otherwise, when the disconnected unit joins the cluster, if `override` is enabled, the cluster renegotiates and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units and any configuration changes made between when the unit was disconnected and reconnected are lost.

Delaying how quickly the primary unit rejoins the cluster when override is enabled

In some cases when `override` is enabled and the unit designated to be the primary unit rejoins the cluster it will become the primary unit too soon and cause traffic disruption. This can happen, for example, if one of the FortiGate interfaces gets its address using PPPoE. If the backup unit is operating as the primary unit and processing traffic, when the primary unit comes up it may need a short time to get a new IP address from the PPPoE server. If the primary unit takes over the cluster before it has an IP address, traffic will be disrupted until the primary unit gets its address.

You can resolve this problem by using the following command to add a wait time. In this example the wait time is 10 seconds. The wait time range is 0 to 3600 seconds and the default wait time is 0 seconds.

```
config system ha
    set override-wait-time 10
end
```

With this wait time configured, after the primary unit is up and running it has 10 seconds to synchronize sessions, get IP address(es) from PPPoE and DHCP servers and so on. After 10 seconds the primary unit sends gratuitous arp packets and all traffic to the cluster is sent to the new primary unit. You can adjust the wait time according to the conditions on your network.

FortiGate HA compatibility with DHCP and PPPoE

FortiGate HA is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. If an interface is

configured for DHCP or PPPoE, turning on high availability may result in the interface receiving an incorrect address or not being able to connect to the DHCP or PPPoE server correctly.



You cannot switch to operate in HA mode if one or more FortiGate interfaces is configured as a PPTP or L2TP client.

You can configure a cluster to act as a DHCP server or a DHCP relay agent. In both active-passive and active-active clusters DHCP relay sessions are always handled by the primary unit. It is possible that a DHCP relay session could be interrupted by a failover. If this occurs the DHCP relay session is not resumed after the failover and the DHCP client may have to repeat the DHCP request.

When a cluster is operating as a DHCP server the primary unit responds to all DHCP requests and maintains the DHCP server address lease database. The cluster also dynamically synchronizes the DHCP server address lease database to the subordinate units. If a failover occurs, the new primary unit will have an up-to-date DHCP server address lease database. Synchronizing the DHCP address lease database prevents the new primary unit from responding incorrectly to new DHCP requests after a failover.

Also, it is possible that when FortiGates first negotiate to form a cluster that a unit that ends up as a subordinate unit in the cluster will have information in its DHCP address lease database that the cluster unit operating as the primary unit does not have. This can happen if a FortiGate responds to DHCP requests while operating as a standalone unit and then when the cluster is formed this unit becomes a subordinate unit. Because of this possibility, after a cluster is formed the DHCP address lease databases of all of the cluster units are merged into one database which is then synchronized to all cluster units.


HA and distributed clustering

The FGCP supports widely separated cluster units installed in different physical locations. Distributed clusters can have cluster units in different rooms in the same building, different buildings in the same location, or even different geographical sites such as different cities, countries or continents.

Just like any cluster, distributed clusters require heartbeat communication between cluster units. In a distributed cluster this heartbeat communication can take place over the Internet or over other transmission methods including satellite linkups.

Most Data Center Interconnect (DCI) or MPLS-based solutions that support layer 2 extensions between the remote data centers should also support HA heartbeat communication between the FortiGates in the distributed locations. Using VLANs and switches in promiscuous mode to pass all traffic between the locations can also be helpful.

HA heartbeat IP addresses are not configurable so the heartbeat interfaces have to be able to communicate over the same subnet. See [HA heartbeat interface IP addresses on page 220](#).



Because of the possible distance it may take a relatively long time for heartbeat packets to be transmitted between cluster units. This could lead to a split brain scenario. To avoid a split brain scenario you can increase the heartbeat interval so that the cluster expects extra time between heartbeat packets. A general rule is to configure the failover time to be longer than the max latency. You could also increase the `hb-lost-threshold` to tolerate losing heartbeat packets if the network connection is less reliable.

In addition you could use different link paths for heartbeat packets to optimize HA heartbeat communication. You could also configure QoS on the links used for HA heartbeat traffic to make sure heartbeat communication has the highest priority.

For information about changing the heartbeat interval and other heartbeat related settings, see [Modifying heartbeat timing on page 222](#).

Clusters of three or four FortiGates

The FGCP supports a cluster of two, three, or four FortiGates. You can add more than two units to a cluster to improve reliability: if two cluster units fail the third will continue to operate and so on. A cluster of three or four units in active-active mode may improve performance since another cluster unit is available for security profile processing. However, active-active FGCP HA results in diminishing performance returns as you add units to the cluster, so the additional performance achieved by adding the third cluster unit may not be worth the cost.

There are no special requirements for clusters of more than two units. Here are a few recommendations though:

- The matching heartbeat interfaces of all of the cluster units must be able to communicate with each other. So each unit's matching heartbeat interface should be connected to the same switch. If the ha1 interface is used for heartbeat communication, then the ha1 interfaces of all of the units in the cluster must be connected together so communication can happen between all of the cluster units over the ha1 interface.

- Redundant heartbeat interfaces are recommended. You can reduce the number of points of failure by connecting each matching set of heartbeat interfaces to a different switch. This is not a requirement; however, and you can connect both heartbeat interfaces of all cluster units to the same switch. However, if that switch fails the cluster will stop forwarding traffic.
- For any cluster, a dedicated switch for each heartbeat interface is recommended because of the large volume of heartbeat traffic and to keep heartbeat traffic off of other networks, but it is not required.
- Full mesh HA can scale to three or four FortiGates. Full mesh HA is not required if you have more than 2 units in a cluster.
- Virtual clustering can only be done with two FortiGates.

Connecting a cluster of three FortiGates

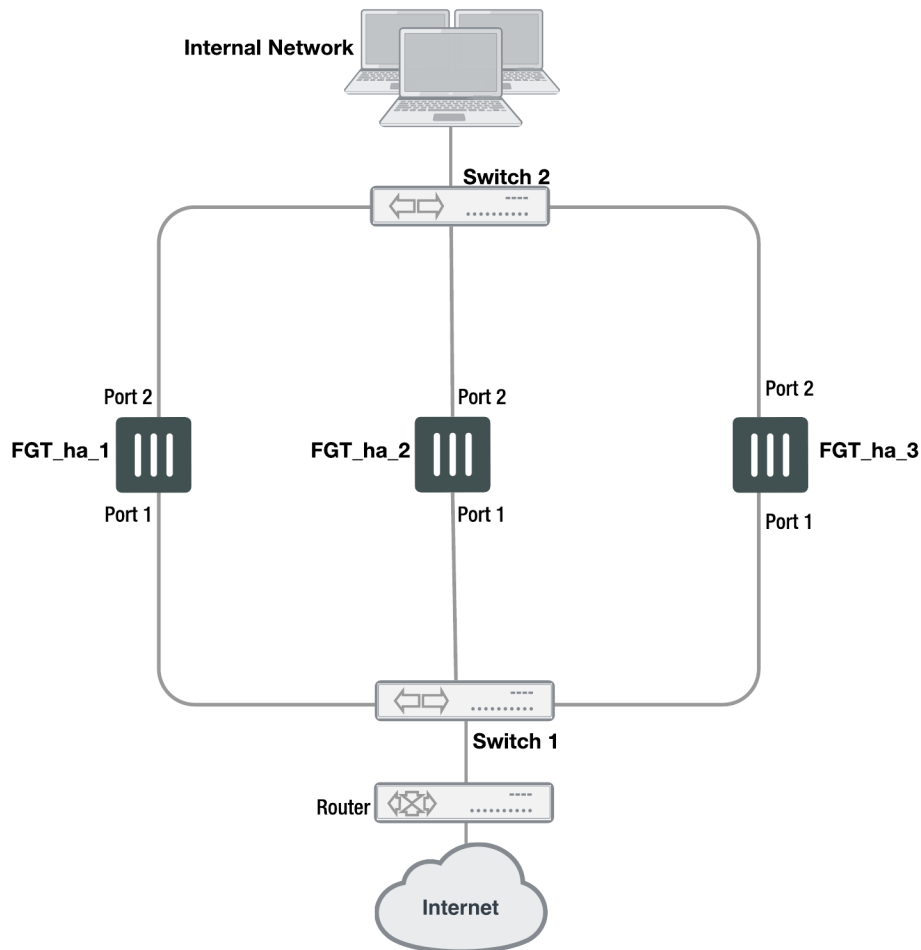
This example shows how to connect a cluster of three FortiGates where:

- Port1 connects the cluster to the Internet
- Port2 connects the cluster to the internal network
- Port3 and Port4 are the heartbeat interfaces

Use the following steps to connect the cluster units to each other and to their networks:

1. Connect the network interfaces:


- Connect the port1 interface of each FortiGate to the same switch (Switch 1) and connect this switch to the Internet.
- Connect the port2 interface of each FortiGate to the same switch (Switch 2) and connect this switch to the internal Network.

Connecting the network interfaces (cluster of three FortiGates)

2. Connect the heartbeat interfaces:

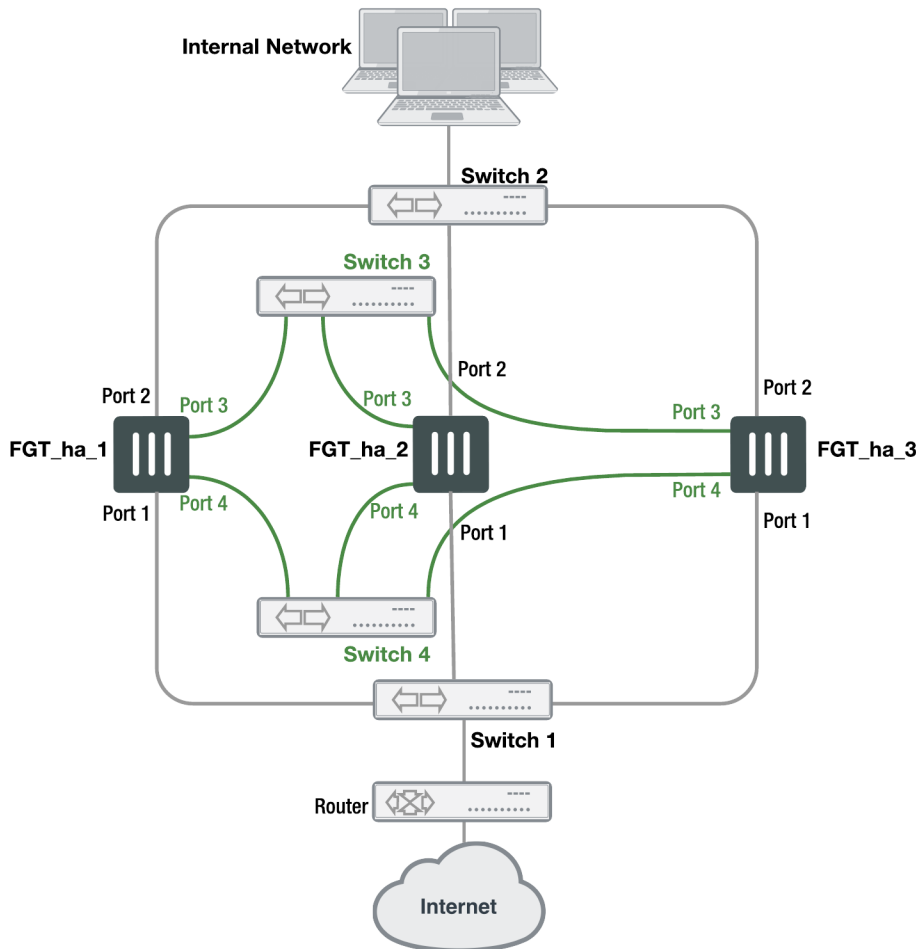
- Connect the port3 interface of each FortiGate to the same switch (Switch 3)
- Connect the port4 interface of each FortiGate to the same switch (Switch 4)

Connecting the heartbeat interfaces (cluster of three FortiGates)



The network and heartbeat connections when combined into one diagram appear like the following:

Network and heartbeat interface connections (cluster of three FortiGates)



Disk storage configuration and HA

If your cluster units include storage disks (for example for storing log messages, WAN optimization data and web caching) all cluster units must have identical storage disk configurations. This means each cluster unit must have same number of disks (including AMC and FortiGate Storage Module (FSM) hard disks) and also means that matching disks in each cluster unit must be the same size, have the same format, and have the same number of partitions.

In most cases the default hard disk configuration of the cluster units will be compatible. However, a hard disk formatted by an older FortiGate firmware version may not be compatible with a hard disk formatted by a more recent firmware version. Problems may also arise if you have used the `execute scsi-dev` command to add or change hard disk protections.

If a cluster unit CLI displays hard disk compatibility messages, you may need to use the `execute scsi-dev delete` command to delete partitions. You can also use the `execute formatlogdisk` command to reformat disks. In some cases after deleting all partitions and reformatting the disks, you may still see disk incompatibility messages. If this happens, contact Fortinet Customer Support for assistance.

FGCP high availability best practices

Fortinet suggests the following practices related to high availability:

- Use Active-Active HA to distribute TCP and UTM sessions among multiple cluster units. An active-active cluster may have higher throughput than a standalone FortiGate unit or than an active-passive cluster.
- Use a different host name on each FortiGate unit when configuring an HA cluster. Fewer steps are required to add host names to each cluster unit before configuring HA and forming a cluster.
- Consider adding an Alias to the interfaces used for the HA heartbeat so that you always get a reminder about what these interfaces are being used for.
- Enabling `load-balance-all` can increase device and network load since more traffic is load-balanced. This may be appropriate for use in a deployment using the firewall capabilities of the FortiGate unit and IPS but no other content inspection.
- An advantage of using session pickup is that non-content inspection sessions will be picked up by the new primary unit after a failover. The disadvantage is that the cluster generates more heartbeat traffic to support session pickup as a larger portion of the session table must be synchronized. Session pickup should be configured only when required and is not recommended for use with SOHO FortiGate models. Session pickup should only be used if the primary heartbeat link is dedicated (otherwise the additional HA heartbeat traffic could affect network performance).
- If session pickup is not selected, after a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Users downloading large files may have to restart their download after a failover. Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.
- If you need to enable session pickup, consider enabling `session-pickup-delay` to improve performance by reducing the number of sessions that are synchronized. See [Session failover \(session-pickup\) on page 262](#).
- Consider using the `session-sync-dev` option to move session synchronization traffic off the HA heartbeat link to one or more dedicated session synchronization interfaces. See [Session failover \(session-pickup\) on page 262](#).
- To avoid unpredictable results, when you connect a switch to multiple redundant or aggregate interfaces in an active-passive cluster you should configure separate redundant or aggregate interfaces on the switch; one for each cluster unit.
- Use SNMP, syslog, or email alerts to monitor a cluster for failover messages. Alert messages about cluster failovers may help find and diagnose network problems quickly and efficiently.

Heartbeat interfaces

Fortinet suggests the following practices related to heartbeat interfaces:



Do not use a FortiGate switch port for the HA heartbeat traffic. This configuration is not supported.

- Configure at least two heartbeat interfaces and set these interfaces to have different priorities.
- For clusters of two FortiGate units, as much as possible, heartbeat interfaces should be directly connected using patch cables (without involving other network equipment such as switches). If switches have to be used they should not be used for other network traffic that could flood the switches and cause heartbeat delays.
 - If you cannot use a dedicated switch, the use of a dedicated VLAN can help limit the broadcast domain to protect the heartbeat traffic and the bandwidth it creates.
- For clusters of three or four FortiGate units, use switches to connect heartbeat interfaces. The corresponding heartbeat interface of each FortiGate unit in the cluster must be connected to the same switch. For improved redundancy use a different switch for each heartbeat interface. In that way if the switch connecting one of the heartbeat interfaces fails or is unplugged, heartbeat traffic can continue on the other heartbeat interfaces and switch.
- Isolate heartbeat interfaces from user networks. Heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. If the cluster consists of two FortiGate units, connect the heartbeat interfaces directly using a crossover cable or a regular Ethernet cable. For clusters with more than two units, connect heartbeat interfaces to a separate switch that is not connected to any network.
- If heartbeat traffic cannot be isolated from user networks, enable heartbeat message encryption and authentication to protect cluster information. See [Enabling or disabling HA heartbeat encryption and authentication on page 223](#).
- Configure and connect redundant heartbeat interfaces so that if one heartbeat interface fails or becomes disconnected, HA heartbeat traffic can continue to be transmitted using the backup heartbeat interface. If heartbeat communication fails, all cluster members will think they are the primary unit resulting in multiple devices on the network with the same IP addresses and MAC addresses (condition referred to as *Split Brain*) and communication will be disrupted until heartbeat communication can be reestablished.
- Do not monitor dedicated heartbeat interfaces; monitor those interfaces whose failure should trigger a device failover.
- Where possible at least one heartbeat interface should not be connected to an NP4 or NP6 processor to avoid NP4 or NP6-related problems from affecting heartbeat traffic.
- Where possible, the heartbeat interfaces should not be connected to an NP4 or NP6 processor that is also processing network traffic.
- Where possible, each heartbeat interface should be connected to a different NP4 or NP6 processor.
- Any FortiGate interface can be used as a heartbeat interface including 10/100/1000Base-T, SFP, QSFP fiber and copper, and so on. If you set up two or more interfaces as heartbeat interfaces each interface can be a different type and speed.

Interface monitoring (port monitoring)

Fortinet suggests the following practices related to interface monitoring (also called port monitoring):

- Wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.
- Monitor interfaces connected to networks that process high priority traffic so that the cluster maintains connections to these networks if a failure occurs.
- Avoid configuring interface monitoring for all interfaces.
- Supplement interface monitoring with remote link failover. Configure remote link failover to maintain packet flow if a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails. See [Remote link failover on page 250](#).

FGCP HA terminology

The following HA-specific terms are used in this document.

Cluster

A group of FortiGates that act as a single virtual FortiGate to maintain connectivity even if one of the FortiGates in the cluster fails.

Cluster unit

A FortiGate operating in a FortiGate HA cluster.

Device failover

Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

Failover

A FortiGate taking over processing network traffic in place of another unit in the cluster that suffered a device failure or a link failure.

Failure

A hardware or software problem that causes a FortiGate or a monitored interface to stop processing network traffic.

FGCP

The FortiGate clustering protocol (FGCP) that specifies how the FortiGates in a cluster communicate to keep the cluster operating.

Full mesh HA

Full mesh HA is a method of removing single points of failure on a network that includes an HA cluster. FortiGate models that support redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA includes redundant connections between all network components. If any single component or any single connection fails, traffic switches to the redundant component or connection.

HA virtual MAC address

When operating in HA mode, all of the interfaces of the primary unit acquire the same HA virtual MAC address. All communications with the cluster must use this MAC address. The HA virtual MAC address is set according to the group ID.

Heartbeat

Also called FGCP heartbeat or HA heartbeat. The heartbeat constantly communicates HA status and synchronization information to make sure that the cluster is operating properly.

Heartbeat device

An Ethernet network interface in a cluster that is used by the FGCP for heartbeat communications among cluster units.

Heartbeat failover

If an interface functioning as the heartbeat device fails, the heartbeat is transferred to another interface also configured as an HA heartbeat device.

Hello state

In the hello state a cluster unit has powered on in HA mode, is using HA heartbeat interfaces to send hello packets, and is listening on its heartbeat interfaces for hello packets from other FortiGates. Hello state may appear in HA log messages.

High availability

The ability that a cluster has to maintain a connection when there is a device or link failure by having another unit in the cluster take over the connection, without any loss of connectivity. To achieve high availability, all FortiGates in the cluster share session and configuration information.

Interface monitoring

You can configure interface monitoring (also called port monitoring) to monitor FortiGate interfaces to verify that the monitored interfaces are functioning properly and connected to their networks. If a monitored interface fails or is disconnected from its network the interface leaves the cluster and a link failover occurs. For more information about interface monitoring, see [Link failover \(port monitoring or interface monitoring\) on page 243](#).

Link failover

Link failover means that if a monitored interface fails, the cluster reorganizes to re-establish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

Load balancing

Also known as active-active HA. All units in the cluster process network traffic. The FGCP employs a technique similar to unicast load balancing. The primary unit interfaces are assigned virtual MAC addresses which are associated on the network with the cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit can process packets itself, or propagate them to subordinate units according to a load balancing schedule. Communication between the cluster units uses the actual cluster unit MAC addresses.

Monitored interface

An interface that is monitored by a cluster to make sure that it is connected and operating correctly. The cluster monitors the connectivity of this interface for all cluster units. If a monitored interface fails or becomes disconnected from its network, the cluster will compensate.

Primary unit

Also called the primary cluster unit, this cluster unit controls how the cluster operates. The primary unit sends hello packets to all cluster units to synchronize session information, synchronize the cluster configuration, and to synchronize the cluster routing table. The hello packets also confirm for the subordinate units that the primary unit is still functioning.

The primary unit also tracks the status of all subordinate units. When you start a management connection to a cluster, you connect to the primary unit.

In an active-passive cluster, the primary unit processes all network traffic. If a subordinate unit fails, the primary unit updates the cluster configuration database.

In an active-active cluster, the primary unit receives all network traffic and re-directs this traffic to subordinate units. If a subordinate unit fails, the primary unit updates the cluster status and redistributes load balanced traffic to other subordinate units in the cluster.

The FortiGate firmware uses the term master to refer to the primary unit.

Session failover

Session failover means that a cluster maintains active network sessions after a device or link failover. FortiGate HA does not support session failover by default. To enable session failover you must change the HA configuration to select Enable Session Pick-up.

Session pickup

If you enable session pickup for a cluster, if the primary unit fails or a subordinate unit in an active-active cluster fails, all communication sessions with the cluster are maintained or picked up by the cluster after the cluster negotiates to select a new primary unit.

If session pickup is not a requirement of your HA installation, you can disable this option to save processing resources and reduce the network bandwidth used by HA session synchronization. In many cases interrupted sessions will resume on their own after a failover even if session pickup is not enabled. You can also enable session pickup delay to reduce the number of sessions that are synchronized by session pickup.

Standby state

A subordinate unit in an active-passive HA cluster operates in the standby state. In a virtual cluster, a subordinate virtual domain also operates in the standby state. The standby state is actually a hot-standby state because the subordinate unit or subordinate virtual domain is not processing traffic but is monitoring the primary unit session table to take the place of the primary unit or primary virtual domain if a failure occurs.

In an active-active cluster all cluster units operate in a work state.

When standby state appears in HA log messages this usually means that a cluster unit has become a subordinate unit in an active-passive cluster or that a virtual domain has become a subordinate virtual domain.

State synchronization

The part of the FGCP that maintains connections after failover.

Subordinate unit

Also called the subordinate cluster unit, each cluster contains one or more cluster units that are not functioning as the primary unit. Subordinate units are always waiting to become the primary unit. If a subordinate unit does not receive hello packets from the primary unit, it attempts to become the primary unit.

In an active-active cluster, subordinate units keep track of cluster connections, keep their configurations and routing tables synchronized with the primary unit, and process network traffic assigned to them by the primary unit. In an active-passive cluster, subordinate units do not process network traffic. However, active-passive subordinate units do keep track of cluster connections and do keep their configurations and routing tables synchronized with the primary unit.

The FortiGate firmware uses the terms slave and subsidiary unit to refer to a subordinate unit.

Virtual clustering

Virtual clustering is an extension of the FGCP for FortiGates operating with multiple VDOMs enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

Work state

The primary unit in an active-passive HA cluster, a primary virtual domain in a virtual cluster, and all cluster units in an active-active cluster operate in the work state. A cluster unit operating in the work state processes traffic, monitors the status of the other cluster units, and tracks the session table of the cluster.

When work state appears in HA log messages this usually means that a cluster unit has become the primary unit or that a virtual domain has become a primary virtual domain.

HA GUI options

Go to **System > HA** to change HA options. You can set the following options to put a FortiGate into HA mode. You can also change any of these options while the cluster is operating.

You can configure HA options for a FortiGate with virtual domains (VDOMs) enabled by logging into the GUI as the global admin administrator and going to **System > HA**.

If already operating in HA mode, go to **System > HA** to display the cluster members list.

Go to **System > HA > View HA Statistics** to view statistics about cluster operation.



If your cluster uses virtual domains, you are configuring HA virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below and see [Virtual clusters on page 143](#).



FortiGate HA is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. If an interface is configured for DHCP or PPPoE, turning on high availability may result in the interface receiving an incorrect address or not being able to connect to the DHCP or PPPoE server correctly.

Mode

Select an HA mode for the cluster or return the FortiGate in the cluster to standalone mode. When configuring a cluster, you must set all members of the HA cluster to the same HA mode. You can select **Standalone** (to disable HA), **Active-Passive**, or **Active-Active**.

If virtual domains are enabled you can select **Active-Passive** or **Standalone**.

Device Priority

Optionally set the device priority of the cluster FortiGate. Each FortiGate in a cluster can have a different device priority. During HA negotiation, the FortiGate with the highest device priority usually becomes the primary unit.

In a virtual cluster configuration, each cluster FortiGate can have two different device priorities, one for each virtual cluster. During HA negotiation, the FortiGate with the highest device priority in a virtual cluster becomes the primary FortiGate for that virtual cluster.

Changes to the device priority are not synchronized. You can accept the default device priority when first configuring a cluster.

Reserve Management Port for Cluster Member

You can provide direct management access to individual cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. See [Managing individual cluster units using a reserved management interface on page 178](#).

Do NOT Synchronize Management VDOM Configuration

This options appears if you have enabled multiple VDOMS and set a VDOM other than the root VDOM to be the management VDOM. You can select this option to prevent the management VDOM configuration from being synchronized between cluster units in the virtual cluster. This allows you to add an interface to the VDOM in each cluster unit and then to give the Interface a different IP address in each cluster unit, allowing you to manage each cluster unit separately.

You can also enable this feature using the following command:

```
config system ha
    set standalone-mgmt-vdom enable
end
```

Group Name

Enter a name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.

Password

Enter a password to identify the cluster. The password must be the same for all cluster FortiGates before the cluster FortiGates can form a cluster.

Two clusters on the same network must have different passwords.

The password is synchronized to all cluster units in an operating cluster. If you change the password of one cluster unit the change is synchronized to all cluster units.

Enable Session pickup

Select to enable session pickup so that if the primary unit fails, sessions are picked up by the cluster unit that becomes the new primary unit.

You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce HA CPU usage and reduce HA heartbeat network bandwidth usage. See [Session failover \(session-pickup\) on page 262](#).

Port Monitor

Select to enable or disable monitoring FortiGate interfaces to verify the monitored interfaces are functioning properly and are connected to their networks. See [Link failover \(port monitoring or interface monitoring\) on page 243](#).

If a monitored interface fails or is disconnected from its network, the interface leaves the cluster and a link failover occurs. The link failover causes the cluster to reroute the traffic being processed by that interface to the same interface of another cluster FortiGate that still has a connection to the network. This other cluster FortiGate becomes the new primary unit.

Port monitoring (also called interface monitoring) is disabled by default. Leave port monitoring disabled until the cluster is operating and then only enable port monitoring for connected interfaces.

You can monitor up to 64 interfaces.

Heartbeat Interface

Select to enable or disable HA heartbeat communication for each interface in the cluster and set the heartbeat interface priority. The heartbeat interface with the highest priority processes all heartbeat traffic. If two or more heartbeat interfaces have the same priority, the heartbeat interface with the lowest hash map order value processes all heartbeat traffic. The GUI lists interfaces in alphanumeric order:

- port1
- port2 through 9
- port10

Hash map order sorts interfaces in the following order:

- port1
- port10
- port2 through port9

The default heartbeat interface configuration is different for each FortiGate model. This default configuration usually sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration or change it as required.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0.

You must select at least one heartbeat interface. If heartbeat communication is interrupted, the cluster stops processing traffic. See [HA heartbeat and communication between cluster units on page 217](#).

You can select up to 8 heartbeat interfaces. This limit only applies to units with more than 8 physical interfaces.

VDOM partitioning

If you are configuring virtual clustering, you can set the virtual domains to be in virtual cluster 1 and the virtual domains to be in virtual cluster 2. The root virtual domain must always be in virtual cluster 1.

FGCP configuration examples and troubleshooting

This chapter contains general procedures and descriptions as well as detailed configuration examples that describe how to configure FortiGate HA clusters.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameter


About the examples in this chapter

The procedures in this chapter describe some of many possible sequences of steps for configuring HA clustering. As you become more experienced with FortiOS HA you may choose to use a different sequence of configuration steps.

For simplicity, many of these procedures assume that you are starting with new FortiGates set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

How to set up FGCP clustering (recommended steps)

This example describes how to enhance the reliability of a network protected by a FortiGate by adding a second FortiGate to create a FortiGate Clustering Protocol (FGCP) High Availability cluster.



The FortiGate already on the network will be configured to become the primary unit by increasing its device priority and enabling override. The new FortiGate will be prepared by setting it to factory defaults to wipe any configuration changes. Then it will be licensed, configured for HA, and then connected to the FortiGate already on the network. All of the FortiGates in a cluster must have the same level of licensing. The new FortiGate becomes the backup unit and its configuration is overwritten by the primary unit.

The recipe contains instructions for both the GUI and the CLI, with some parts of the configuration requiring use of the CLI.

Before you start the FortiGates should be running the same FortiOS firmware version and interfaces should not be configured to get their addresses from DHCP or PPPoE.

1. Configuring the primary FortiGate

If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license.

Connect to the primary FortiGate and go to **System > Dashboard > Status** and locate the **System Information** widget. Change the unit's **Host Name** to identify it as the primary FortiGate.

Current Name FG100D3G12804410

New Name

You can also enter this CLI command:

```
config system global
    set hostname Primary_FortiGate
end
```

If you have not already done so, register the primary FortiGate and apply licenses to it before setting up the cluster. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMs). All of the FortiGates in a cluster must have the same level of licensing. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

License Information

Support Contract	Registration	Registered (bdickie@fortinet.com)	Launch Portal
FortiGuard	IPS & Application Control	Licensed (Expires 2016-08-22)	
	AntiVirus	Licensed (Expires 2016-08-22)	
	Web Filtering	Licensed (Expires 2016-08-21)	
	Anti-Spam Filtering	Licensed (Expires 2016-08-21)	
FortiCloud	Account		Activate
FortiSandbox	FortiSandbox Appliance	Not Configured	Configure
FortiClient	Status	Free License	How to Purchase
	Clients Registered	0 of 10	Enter License
	FortiClient Installers		Details
FortiToken Mobile	Tokens Assigned	0 of 2	

Enter this CLI command to set the HA mode to active-passive, set a group name and password, increase the device priority to a higher value (for example, 250) and enable override.

```
config system ha
  set mode a-p
  set group-name My-HA-Cluster
  set password
  set priority 250
  set override enable
  set hbdev ha1 50 ha2 50
end
```

Enabling override and increasing the device priority means this unit should always become the primary unit.

This command also selects ha1 and ha2 to be the heartbeat interfaces and sets their priorities to 50.

You can also use the GUI to configure most of these settings.

Mode: Active-Passive

Device Priority: 250

☐ Reserve Management Port for Cluster Member Internal

Cluster Settings

Group Name: My-HA-Cluster

Password: *****

☐ Enable Session Pick-up

	Port Monitor	Heartbeat Interface	
		Enable	Priority(0-512)
dmz (DMZ server network)	<input type="checkbox"/>	<input type="checkbox"/>	0
ha1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
ha2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
mgmt	<input type="checkbox"/>		
port9	<input type="checkbox"/>	<input type="checkbox"/>	0
port10	<input type="checkbox"/>	<input type="checkbox"/>	0
port11	<input type="checkbox"/>	<input type="checkbox"/>	0
port14	<input type="checkbox"/>	<input type="checkbox"/>	0
port15	<input type="checkbox"/>	<input type="checkbox"/>	0
port16	<input type="checkbox"/>	<input type="checkbox"/>	0
wan1	<input type="checkbox"/>	<input type="checkbox"/>	0
wan2	<input type="checkbox"/>	<input type="checkbox"/>	0

Override can only be enabled from the CLI.

```
config system ha
  set override enable
end
```

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as FGCP negotiation takes place and the MAC addresses of the FortiGate are changed to HA virtual MAC addresses. These virtual MAC addresses are used for failover. The actual virtual MAC address assigned to each FortiGate interface depends on the HA group ID. Since this example does not involve changing the HA group ID, the FortiGate's interfaces will have the following MAC addresses: 00:09:0f:09:00:00, 00:09:0f:09:00:01, 00:09:0f:09:00:02 and so on.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You can usually delete the arp table from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) command to view the virtual MAC address of any FortiGate interface. Depending on the FortiGate model, the output from this command could include lines similar to the following:

```
Current_HWaddr: 00:09:0f:09:00:00
Permanent_HWaddr 02:09:0f:78:18:c9
```

2. Configuring the backup FortiGate

Enter this command to reset the new FortiGate to factory default settings.

```
execute factoryreset
```

You can skip this step if the new FortiGate is fresh from the factory. But if its configuration has been changed at all it is recommended to set it back to factory defaults to reduce the chance of synchronization problems.

Change the firmware running on the new FortiGate to be the same version as is running on the primary unit.

Register the backup FortiGate and apply licenses to it before adding it to the cluster. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMs).

License Information

Support Contract	Registration	Registered (bdickie@fortinet.com)	Launch Portal
FortiGuard	IPS & Application Control	Licensed (Expires 2016-08-22)	
	AntiVirus	Licensed (Expires 2016-08-22)	
	Web Filtering	Licensed (Expires 2016-08-21)	
	Anti-Spam Filtering	Licensed (Expires 2016-08-21)	
FortiCloud	Account		Activate
FortiSandbox	FortiSandbox Appliance	Not Configured	Configure
FortiClient	Status	Free License	How to Purchase
	Clients Registered	0 of 10	Enter License
	FortiClient Installers		Details
FortiToken Mobile	Tokens Assigned	0 of 2	

Go to **System > Dashboard**
> Status and change the unit's
Host Name to identify it as the
 backup FortiGate.

Current Name FG100D3G12801361

New Name Backup_FortiGate

You can also enter this CLI
 command:

```
config system global
    set hostname Backup_FortiGate
end
```

Duplicate the primary unit HA settings, except set the device priority to a lower value and do not enable override.

You can configure all of these settings from the GUI.

Mode Active-Passive

Device Priority 50

☐ Reserve Management Port for Cluster Member Internal

Cluster Settings

Group Name My-HA-Cluster

Password *****

☐ Enable Session Pick-up

	Port Monitor	Heartbeat Interface	
		Enable	Priority(0-512)
dmz (DMZ server network)	<input type="checkbox"/>	<input type="checkbox"/>	0
ha1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
ha2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
mgmt	<input type="checkbox"/>		
port9	<input type="checkbox"/>	<input type="checkbox"/>	0
port10	<input type="checkbox"/>	<input type="checkbox"/>	0
port11	<input type="checkbox"/>	<input type="checkbox"/>	0
port14	<input type="checkbox"/>	<input type="checkbox"/>	0
port15	<input type="checkbox"/>	<input type="checkbox"/>	0
port16	<input type="checkbox"/>	<input type="checkbox"/>	0
wan1	<input type="checkbox"/>	<input type="checkbox"/>	0
wan2	<input type="checkbox"/>	<input type="checkbox"/>	0

You can also enter this CLI command:

```
config system ha
  set mode a-p
  set group-name My-HA-Cluster
  set password
  set priority 50
  set hbdev ha1 50 ha2 50
end
```

3. Connecting the cluster

Connect the HA cluster as shown in the initial diagram. Making these connections will disrupt network traffic as you disconnect and re-connect cables.

When connected the primary and backup FortiGates find each other and negotiate to form an HA cluster. The Primary unit synchronizes its configuration with the backup FortiGate. Forming the cluster happens automatically with minimal or no disruption to network traffic.

4. Checking cluster operation and disabling override

Check the cluster synchronization status to make sure the primary and backup units have the same configuration. Log into the primary unit CLI and enter this command:

```
diag sys ha cluster-csum
```

The CLI lists all members' checksums. If both cluster units have identical checksums you can be sure that their configurations are synchronized. If the checksums are different wait a short while and enter the command again. Repeat until the checksums are identical. It may take a while for some parts of the configuration to be synchronized. If the checksums never become identical contact Fortinet support to help troubleshoot the problem.



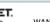
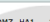
Disable override on the primary unit (recommended).

```
config system ha
    set override disable
end
```

The HA cluster dynamically responds to network conditions. If you keep override enabled the same FortiGate will always be the primary FortiGate. Because of this, however; the cluster may negotiate more often potentially disrupting traffic.

If you disable override it is more likely that the new FortiGate could become the primary unit. Disabling override is recommended unless its important that the same FortiGate remains the primary unit.

Connect to the primary FortiGate GUI and go to **System > HA** to view the cluster information.

Cluster Member		Hostname	Serial No.	Role	Priority
 		Primary_FortiGate	FG100D3G12804410	MASTER	128
		Backup_FortiGate	FG100D3G12801361	SLAVE	50

Select **View HA Statistics** for more information on how the cluster is operating and processing traffic.

Unit	Status	Up Time	Monitor			
Primary_FortiGate FG100D3G12804410	✓	0 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		1 hours	1%	26	81857	0
		44 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		2 seconds	34%	78 Kbps	27300058	0
Backup_FortiGate FG100D3G12801361	✓	2 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		0 hours	0%	6	8718576	0
		15 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		15 seconds	19%	13 Kbps	2778691497	0

5. Results

Normally, traffic should now be flowing through the primary FortiGate. However, if the primary FortiGate is unavailable, traffic should failover and the backup FortiGate will be used. Failover will also cause the primary and backup FortiGates to reverse roles, even when both FortiGates are available again.

To test this, ping the IP address 8.8.8.8 using a PC on the internal network. After a moment, power off the primary FortiGate.

```

Reply from 8.8.8.8: bytes=32 time=50ms TTL=53
Reply from 8.8.8.8: bytes=32 time=38ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Request timed out.
Reply from 8.8.8.8: bytes=32 time=482ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Reply from 8.8.8.8: bytes=32 time=36ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Reply from 8.8.8.8: bytes=32 time=38ms TTL=53

```

You will see a momentary pause in the Ping results, until traffic diverts to the backup FortiGate, allowing the Ping traffic to continue.

If you are using port monitoring, you can also unplug the primary FortiGate's Internet-facing interface to test failover.


Setting up two new FortiGates as an FGCP cluster

This section describes a simple HA network topology that includes an HA cluster of two FortiGates in NAT/Route mode installed between an internal network and the Internet. The example uses a generic FortiGate with four interfaces named port1, port2, port3 and port4.

Example NAT/Route mode HA network topology

The figure below shows a typical FortiGate HA cluster consisting of two FortiGates (FGT_ha_1 and FGT_ha_2) connected to the same internal (port2) and external (port1) networks.

Example NAT/Route mode HA network topology



Port3 and port4 are used as the heartbeat interfaces. Because the cluster consists of two FortiGates, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

General configuration steps

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. Confirm that the cluster units are operating as a cluster and add basic configuration settings to the cluster.
 - View cluster status from the GUI or CLI.
 - Add a password for the admin administrative account.

- Change the IP addresses and netmasks of the internal and external interfaces.
- Add a default route.

Configuring a NAT/Route mode active-passive cluster of two FortiGates - GUI

Use the following procedures to configure two FortiGates for NAT/Route HA operation using the GUI. These procedures assume you are starting with two FortiGates with factory default settings.



Give each cluster unit a unique host name to make the individual units easier to identify when they are part of a functioning cluster. The default host name is the FortiGate serial number. You may want to change this host name to something more meaningful for your network.

To configure the first FortiGate (host name FGT_ha_1)

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	<div> ✓ Registered (bdickie@fortinet.com) Launch Portal </div>
	FortiGuard	IPS & Application Control	<div> ✓ Licensed (Expires 2016-08-22) </div>
		AntiVirus	<div> ✓ Licensed (Expires 2016-08-22) </div>
		Web Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
		Anti-Spam Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
	FortiCloud	Account	<div> Activate </div>
	FortiSandbox	FortiSandbox Appliance	<div> ✗ Not Configured Configure </div>
	FortiClient	Status	<div> ✓ Free License How to Purchase </div>
		Clients Registered	<div> <input type="text"/> 0 of 10 Enter License </div>
		FortiClient Installers	<div> <input type="text"/> Details </div>
	FortiToken Mobile	Tokens Assigned	<div> <input type="text"/> 0 of 2 </div>

2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. On the **System Information** dashboard widget beside **Host Name**, select **Change**.
4. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

5. Select **OK**.
6. Go to **System > HA** and change the following settings:

Mode	Active-Passive
Group Name	example1.com
Password	HA_pass_1



This is the minimum recommended configuration for an active-passive HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each unit in the cluster.

7. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HWaddr:  00:09:0f:09:00:00
Permanent_HWaddr 02:09:0f:78:18:c9
.
.
.
```

10. Power off the first FortiGate (FGT_ha_1).



Note the details and format of the output of the `get hardware nic` command are specific to the interface hardware. Different FortiGate models and even different interfaces in the same FortiGate may have different output.

To configure the second FortiGate (host name FGT_ha_2)

1. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
2. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_2
-----------------	----------

3. Select **OK**.
4. Go to **System > HA** and change the following settings:

Mode	Active-Passive
Group Name	example1.com
Password	HA_pass_1

5. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

6. Power off the second FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate. You could have performed the following configuration steps separately on each FortiGate before you connected them to form a cluster.

1. Start Internet Explorer and browse to the address <https://192.168.1.99> (remember to include the “s” in https://). The FortiGate Login is displayed.
2. Type `admin` in the **Name** field and select Login.
The FortiGate dashboard is displayed.
The **System Information** dashboard widget shows the **Cluster Name** (example1.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
3. Go to **System > HA** to view the cluster members list.
The list shows both cluster units, their host names, their roles in the cluster, and their device priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster to connect to its network. The following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster GUI.
2. Go to **System > Admin > Administrators**.
3. Edit `admin` and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **System > Network > Interfaces**.
7. Edit the `port2` interface and change **IP/Netmask** to 10.11.101.100/24.
8. Select **OK**.



After changing the IP address of the port1 interface you may have to change the IP address of your management computer and then reconnect to the port1 interface using the 172.20.120.141 IP address.

9. Edit the `port1` interface and change **IP/Netmask** to 172.20.120.141/24.
10. Select **OK**.
11. Go to **Router > Static > Static Routes**.
12. Change the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	port1
Distance	10

13. Select **OK**.

Configuring a NAT/Route mode active-passive cluster of two FortiGates - CLI

Use the following procedures to configure two FortiGates for NAT/Route HA operation using the FortiGate CLI. These procedures assume you are starting with two FortiGates with factory default settings.

To configure the first FortiGate (host name FGT_ha_1)

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal (or any terminal emulation program), enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.
The FortiGate CLI login prompt appears.
If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.
7. Type `admin` and press **Enter** twice.
8. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS).
9. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
FortiToken licenses can be added at any time because they are synchronized to all cluster members.

10. Change the host name for this FortiGate.

```
config system global
  set hostname FGT_ha_1
end
```

11. Configure HA settings.

```
config system ha
  set mode a-p
  set group-name example1.com
  set password HA_pass_1
end
```

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the

FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HAaddr      00:09:0f:09:00:00
Permanent_HWaddr    02:09:0f:78:18:c9
.
.
.
```

10. Display the HA configuration (optional).

```
get system ha
  group-id : 0
  group-name : example1.com
  mode : a-p
  password : *
  hbdev : "port3" 50 "port4" 50
  session-sync-dev :
  route-ttl : 10
  route-wait : 0
  route-hold : 10
  sync-config : enable
  encryption : disable
  authentication : disable
  hb-interval : 2
  hb-lost-threshold : 20
  hello-holddown : 20
  arps : 5
  arps-interval : 8
  session-pickup : disable
  update-all-session-timer: disable
  session-sync-daemon-number: 1
  link-failed-signal : disable
  uninterruptible-upgrade: enable
  ha-mgmt-status : disable
  ha-eth-type : 8890
  hc-eth-type : 8891
  l2ep-eth-type : 8893
  ha-uptime-diff-margin: 300
```



```

vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
slave-switch-standby: disable
minimum-worker-threshold: 1
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"

```

11. Power off the FortiGate.

To configure the second FortiGate (host name FGT_ha_2)

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.
The FortiGate CLI login prompt appears.
7. Type `admin` and press **Enter** twice.
8. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS).
9. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
FortiToken licenses can be added at any time because they are synchronized to all cluster members.

10. Change the host name for this FortiGate.

```

config system global
    set hostname FGT_ha_2
end

```

11. Configure HA settings.

```

config system ha
    set mode a-p
    set group-name example1.com

```

```
set password HA_pass_1
end
```

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

12. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example1.com
mode : a-p
password : *
hbdev : "port3" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
hello-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
update-all-session-timer: disable
session-sync-daemon-number: 1
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
ha-uptime-diff-margin: 300
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
slave-switch-standby: disable
minimum-worker-threshold: 1
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"
```

13. Power off the FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Determine which cluster unit is the primary unit.
 - Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
 - Enter the command `get system status`.
 - If the command output includes `Current HA mode: a-p, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step ["Setting up two new FortiGates as an FGCP cluster" on page 73](#).
 - If the command output includes `Current HA mode: a-p, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

3. Check the cluster synchronization status to make sure the primary and backup units have the same configuration. Log into the primary unit CLI and enter this command:

```
diag sys ha cluster-csum
```

The CLI lists all members' checksums. If both cluster units have identical checksums you can be sure that their configurations are synchronized. If the checksums are different wait a short while and enter the command again. Repeat until the checksums are identical. It may take a while for some parts of the configuration to be synchronized. If the checksums never become identical contact Fortinet support to help troubleshoot the problem.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to add some basic settings to the cluster so that it can connect to the network.

1. Log into the primary unit CLI.
2. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <password_str>
  end
```

3. Configure the port1 and port2 interfaces.

```
config system interface
  edit port1
    set ip 172.20.120.141/24
  next
  edit port2
    set ip 10.11.101.100/24
  end
```

4. Add a default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device port1
  end
```

Adding a new FortiGate to an operating cluster

This procedure describes how to add a new FortiGate to a functioning cluster. Adding a new unit to a cluster does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the new cluster unit.

You can use this procedure to add as many units as required to the cluster.

To add a new unit to a functioning cluster

1. Install the same firmware build on the new cluster unit as is running on the cluster.
2. Register and apply licenses to the new cluster unit. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
3. If you are planning on adding FortiToken licenses you can do that now and configure FortiTokens or you can wait until you have formed the cluster and then add the FortiToken licenses and tokens.
4. Configure the new cluster unit for HA operation with the same HA configuration as the other units in the cluster.
5. If the cluster is running in Transparent mode, change the operating mode of the new cluster unit to Transparent mode.
6. Power off the new cluster unit.
7. Connect the new cluster unit to the cluster.
8. For example, see [How to set up FGCP clustering \(recommended steps\) on page 66](#).
9. Power on the new cluster unit.

When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the new unit configuration with the configuration of the primary unit.

You can add a new unit to a functioning cluster at any time. For best results the new cluster unit should:

- Have the same hardware version as the cluster units.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.

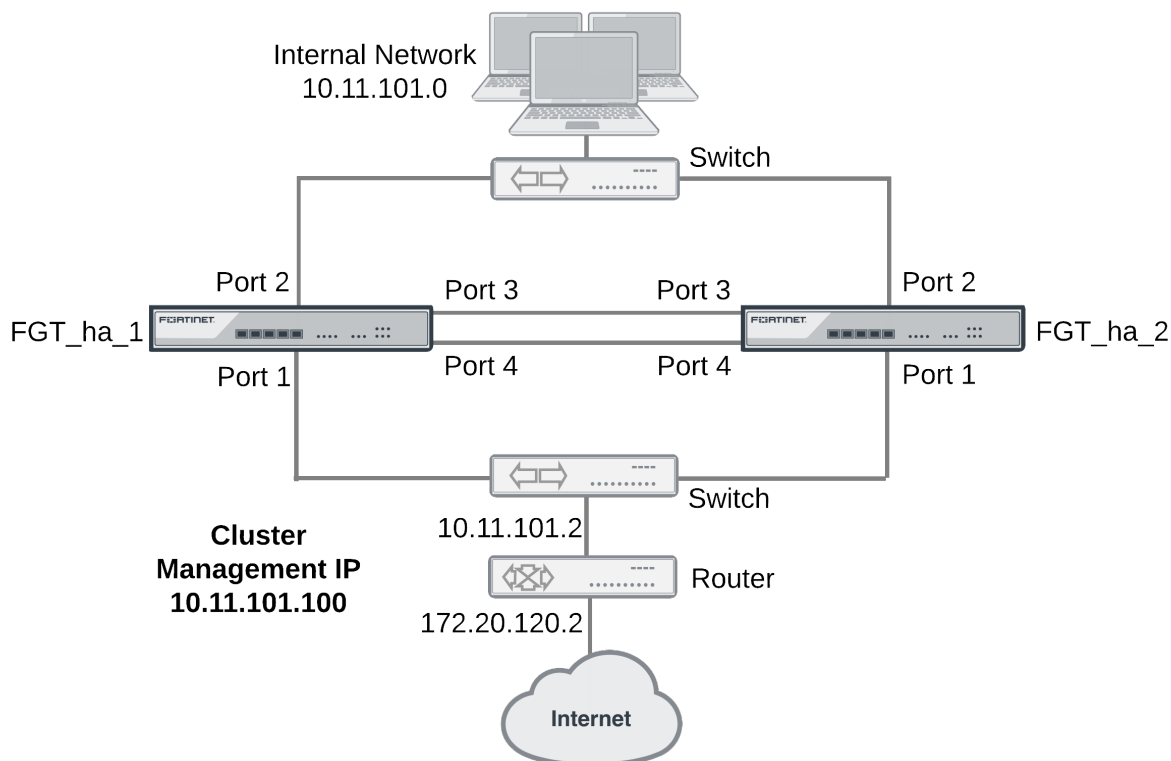
Active-active HA cluster in Transparent mode

This section describes a simple HA network topology that includes an HA cluster of two generic FortiGates installed between an internal network and the Internet and running in Transparent mode.

Example Transparent mode HA network topology

The figure below shows a Transparent mode FortiGate HA cluster consisting of two FortiGates (FGT_ha_1 and FGT_ha_2) installed between the Internet and internal network. The topology includes a router that performs NAT between the internal network and the Internet. The cluster management IP address is 10.11.101.100.

Transparent mode HA network topology



Port3 and port4 are used as the heartbeat interfaces. Because the cluster consists of two FortiGates, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

General configuration steps

This section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

In this example, the configuration steps are identical to the NAT/Route mode configuration steps until the cluster is operating. When the cluster is operating, you can switch to Transparent mode and add basic configuration settings to cluster.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. Confirm that the cluster units are operating as a cluster.
4. Switch the cluster to Transparent mode and add basic configuration settings to the cluster.

- Switch to Transparent mode, add the management IP address and a default route.
- Add a password for the admin administrative account.
- View cluster status from the GUI or CLI.

Configuring a Transparent mode active-active cluster of two FortiGates - GUI

Use the following procedures to configure the FortiGates for HA operation using the FortiGate GUI. These procedures assume you are starting with two FortiGates with factory default settings.



Waiting until you have established the cluster to switch to Transparent mode means fewer configuration steps because you can switch the mode of the cluster in one step.

To configure the first FortiGate (host name FGT_ha_1)

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	<div> ✓ Registered (bdickie@fortinet.com) Launch Portal </div>
	FortiGuard	IPS & Application Control	<div> ✓ Licensed (Expires 2016-08-22) </div>
		AntiVirus	<div> ✓ Licensed (Expires 2016-08-22) </div>
		Web Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
		Anti-Spam Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
	FortiCloud	Account	<div> Activate </div>
	FortiSandbox	FortiSandbox Appliance	<div> ✗ Not Configured Configure </div>
	FortiClient	Status	<div> ✓ Free License <div> How to Purchase Enter License </div> </div>
		Clients Registered	<div> <input type="text"/> 0 of 10 Details </div>
		FortiClient Installers	
	FortiToken Mobile	Tokens Assigned	<div> <input type="text"/> 0 of 2 </div>

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

2. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

4. Select **OK**.
5. Go to **System > HA** and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

6. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select **OK** you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HAaddr    00:09:0f:09:00:00
Permanent_HWaddr 02:09:0f:78:18:c9
.
.
.
```

10. Power off the first FortiGate.

To configure the second FortiGate (host name FGT_ha_2)

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	<div> ✓ Registered (bdickie@fortinet.com) Launch Portal </div>
	FortiGuard	IPS & Application Control	<div> ✓ Licensed (Expires 2016-08-22) </div>
		AntiVirus	<div> ✓ Licensed (Expires 2016-08-22) </div>
		Web Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
		Anti-Spam Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
	FortiCloud	Account	<div> ⚙️ Activate </div>
	FortiSandbox	FortiSandbox Appliance	<div> ✗ Not Configured Configure </div>
	FortiClient	Status	<div> ✓ Free License <div> How to Purchase Enter License </div> </div>
		Clients Registered	<div> <input type="text"/> 0 of 10 Details </div>
		FortiClient Installers	
	FortiToken Mobile	Tokens Assigned	<div> <input type="text"/> 0 of 2 </div>

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

- On the **System Information** dashboard widget, beside **Host Name** select **Change**.
- Enter a new Host Name for this FortiGate.

New Name	FGT_ha_2
-----------------	----------

- Select **OK**.
- Go to **System > HA** and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2

- Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

- Power off the second FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To switch the cluster to Transparent mode

Switching from NAT/Route to Transparent mode involves adding the Transparent mode management IP address and default route.



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

1. Start a web browser and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).
- The FortiGate Login is displayed.
2. Type admin in the Name field and select Login.
3. Under System Information, beside **Operation Mode** select **Change**.
4. Set Operation Mode to Transparent.
5. Configure basic Transparent mode settings.

Operation Mode	Transparent
Management IP/Mask	10.11.101.100/24
Default Gateway	10.11.101.2

6. Select **Apply**.

The cluster switches to operating in Transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate. You could have performed the following configuration steps separately on each FortiGate before you connected them to form a cluster.

1. Start Internet Explorer and browse to the address <https://10.11.101.100> (remember to include the “s” in https://). The FortiGate Login is displayed.
2. Type admin in the Name field and select Login. The FortiGate dashboard is displayed.

The System Information dashboard widget shows the **Cluster Name** (example2.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.

3. Go to **System > HA** to view the cluster members list. The list shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster. Note that the following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster GUI.
2. Go to **System > Admin > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.



You added a default gateway when you switched to Transparent mode so you don't need to add a default route as part of the basic configuration of the cluster at this point.

Configuring a Transparent mode active-active cluster of two FortiGates - CLI

Use the following procedures to configure the FortiGates for Transparent mode HA operation using the FortiGate CLI.

To configure each FortiGate for HA operation

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.

3. Start HyperTerminal, enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.
The FortiGate CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.
7. Type `admin` and press **Enter** twice.
8. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
9. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
10. Change the host name for this FortiGate. For example:

```
config system global
    set hostname FGT_ha_1
end
```

11. Configure HA settings.

```
config system ha
    set mode a-a
    set group-name example2.com
    set password HA_pass_2
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HAaddr      00:09:0f:09:00:00
Permanent_HWaddr    02:09:0f:78:18:c9
.
.
.
```

10. Display the HA configuration (optional).

```
get system ha
  group-id : 0
  group-name : example2.com
  mode : a-a
  password : *
  hbdev : "port3" 50 "port4" 50
  session-sync-dev :
  route-ttl : 10
  route-wait : 0
  route-hold : 10
  sync-config : enable
  encryption : disable
  authentication : disable
  hb-interval : 2
  hb-lost-threshold : 20
  hello-holddown : 20
  arps : 5
  arps-interval : 8
  session-pickup : disable
  update-all-session-timer: disable
  session-sync-daemon-number: 1
  link-failed-signal : disable
  uninterruptible-upgrade: enable
  ha-mgmt-status : disable
  ha-eth-type : 8890
  hc-eth-type : 8891
  l2ep-eth-type : 8893
  ha-uptime-diff-margin: 300
  vcluster2 : disable
  vcluster-id : 1
  override : disable
```

```

priority : 128
slave-switch-standby: disable
minimum-worker-threshold: 1
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"

```

11. Power off the FortiGate.

To configure the second FortiGate (host name FGT_ha_2)

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.
The FortiGate CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.
7. Type `admin` and press **Enter** twice.
8. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
9. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
10. Change the host name for this FortiGate.


```

config system global
    set hostname FGT_ha_2
end

```
11. Configure HA settings.


```

config system ha

```

```
set mode a-a
set group-name example2.com
set password HA_pass_2
end
```

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

12. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example2.com
mode : a-a
password : *
hbdev : "port3" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
hello-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
update-all-session-timer: disable
session-sync-daemon-number: 1
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
ha-uptime-diff-margin: 300
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
schedule : round-robin
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"
schedule : round-robin
```

13. Power off the FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To connect to the cluster CLI and switch the cluster to Transparent mode

1. Determine which cluster unit is the primary unit.
 - Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
 - Enter the command `get system status`.
 - If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.
 - If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode.

2. Change to transparent mode.

```
config system settings
  set opmode transparent
  set manageip 192.168.20.3/24
  set gateway 192.168.20.1
end
```

The cluster switches to Transparent Mode, and your administration session is disconnected.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (192.168.20.3).

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Determine which cluster unit is the primary unit.
 - Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
 - Enter the command `get system status`.
 - If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step ["Active-active HA cluster in Transparent mode" on page 85](#).

- If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

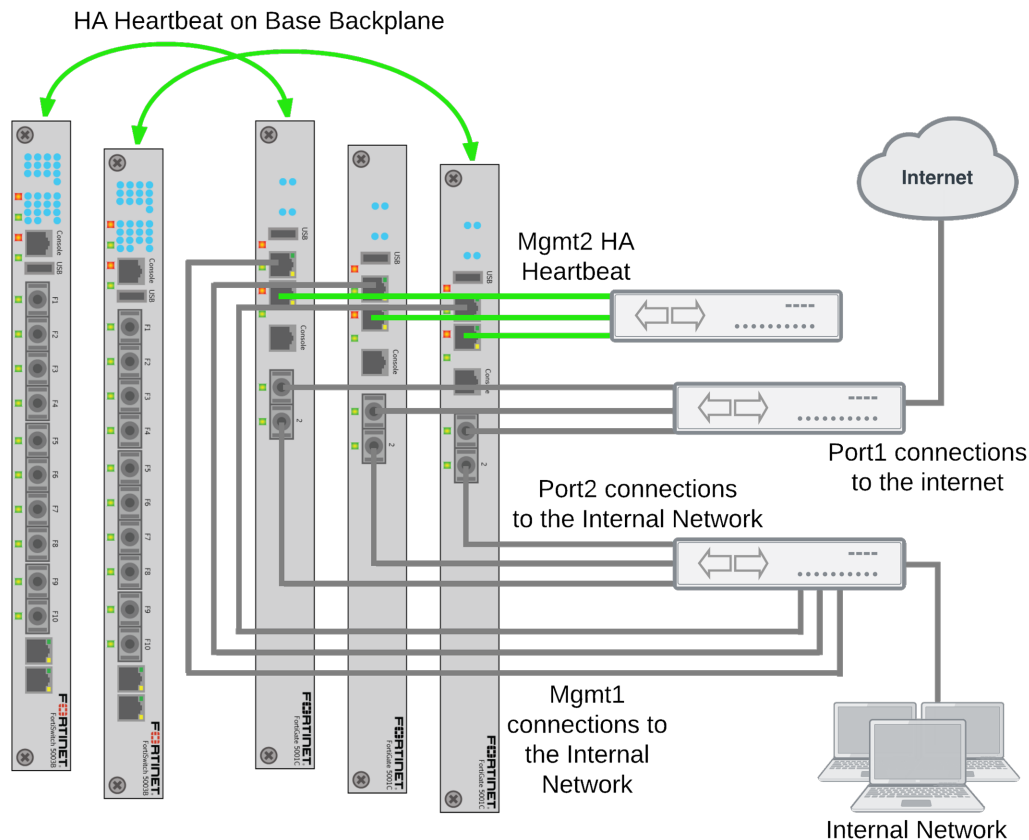
FortiGate-5000 active-active HA cluster with FortiClient licenses

This section describes how to configure an HA cluster of three FortiGate-5001D units that connect an internal network to the Internet. The FortiGate-5001D units each have a FortiClient license installed on them to support FortiClient profiles.

Normally it is recommended that you add FortiClient licenses to the FortiGates before setting up the cluster. This example; however, describes how to apply FortiClient licenses to the FortiGates in an operating cluster.

Example network topology

The following diagram shows an HA cluster consisting of three FortiGate-5001D cluster units (host names slot-3, slot-4, and slot-5) installed in a FortiGate-5000 series chassis with two FortiController-5003B units for heartbeat communication between the cluster units. The cluster applies security features including FortiClient profiles to data traffic passing through it.



The cluster is managed from the internal network using the FortiGate-5001C mgmt1 interfaces configured as HA reserved management interfaces. Using these reserved management interfaces the overall cluster can be managed and cluster units can be managed individually. Individual management access to each cluster unit makes some operations, such as installing FortiClient licenses, easier and also allows you to view status of each cluster unit.

The reserved management interface of each cluster unit has a different IP address and retains its own MAC address. The cluster does not change the reserved management interface MAC address.

Example network topology

By default base1 and base2 are used for heartbeat communication between the FortiGates. To use the base1 and base2 interfaces for the HA heartbeat, the example describes how to display the backplane interfaces on the GUI before turning on HA.

This example also includes using the mgmt2 interface for heartbeat communication for additional heartbeat redundancy.

To connect the cluster

1. Connect the FortiGate-5001C port1 interfaces to a switch and connect that switch to the Internet.
2. Connect the FortiGate-5001C port2 interfaces to a switch and connect that switch to the internal network.
3. Connect the FortiGate-5001C mgmt1 interfaces to a switch that connects to the engineering network.
4. Connect the FortiGate-5001C mgmt2 interfaces to a switch for heartbeat communication between them.

Configuring the FortiGate-5000 active-active cluster - GUI

These procedures assume you are starting with three FortiGate-5001C boards and two FortiSwitch-5003B boards installed in a compatible FortiGate-5000 series chassis. The FortiSwitch-5003B boards are in chassis slots 1 and 2 and the FortiGate-5001C boards are in chassis slots 3, 4, and 5 and the chassis is powered on. All devices are in their factory default configuration. No configuration changes to the FortiSwitch-5003B boards are required.

To configure the FortiGate-5001C units

1. From the internal network, log into the GUI of the FortiGate-5001C unit in chassis slot 3 by connecting to the mgmt1 interface.



By default the mgmt1 interface of each FortiGate-5001C unit has the same IP address. To log into each FortiGate-5001C unit separately you could either disconnect the mgmt1 interfaces of the units that you don't want to log into or change the mgmt1 interface IP addresses for each unit by connecting to each unit's CLI from their console port.

2. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). In this example you will leave **FortiClient** licensing until after you have formed the cluster. Otherwise, all of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	<div> ✓ Registered (bdickie@fortinet.com) Launch Portal </div>
	FortiGuard	IPS & Application Control	<div> ✓ Licensed (Expires 2016-08-22) </div>
		AntiVirus	<div> ✓ Licensed (Expires 2016-08-22) </div>
		Web Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
		Anti-Spam Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
	FortiCloud	Account	<div> Activate </div>
	FortiSandbox	FortiSandbox Appliance	<div> ✗ Not Configured Configure </div>
	FortiClient	Status	<div> ✓ Free License <div> How to Purchase Enter License </div> </div>
		Clients Registered	<div> <input type="text"/> 0 of 10 Details </div>
		FortiClient Installers	
	FortiToken Mobile	Tokens Assigned	<div> <input type="text"/> 0 of 2 </div>

- You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
- On the **System Information** dashboard widget, beside **Host Name** select **Change**.
- Enter a new Host Name for this FortiGate, for example:

New Name	5001C-Slot-3
-----------------	--------------

- Connect to the CLI and enter the following command to display backplane interfaces on the GUI:

```
config system global
    set show-backplane-intf enable
end
```

- Set the Administrative Status of the base1 and base 2 interfaces to **Up**.
You can do this from the GUI by going to **Network > Interfaces**, editing each interface and setting **Administrative Status** to **Up**.

You can also do this from the CLI using the following command:

```
config system interface
    edit base1
        set status up
    next
    edit base2
        set status up
    end
```

- Go to **Network > Interfaces** and configure the IP address of the mgmt1 interface.

Because mgmt1 will become the reserved management interface for the cluster unit each FortiGate-5001C should have a different mgmt1 interface IP address. Give the mgmt1 interface an address that is valid for the internal network. Once HA with the reserved Management interface is enabled the IP address of the mgmt1 interface can be on the same subnet as the port2 interface (which will also be connected to the Internal network).

After the FortiGate is operating in HA mode the mgmt1 interface will retain its original MAC address instead of being assigned a virtual MAC address.

9. Go to **System > HA** and change the following settings:

Set the **Mode** to **Active-Active**.

Select **Reserve Management Port for Cluster Member** and select **mgmt1**.

Set the group name and password:

Group Name	example3.com
Password	HA_pass_3

Set the Heartbeat interface configuration to use base1, base2 and mgmt2 for heartbeat communication. Set the priority of each heartbeat interface to 50:

Heartbeat Interface		
	Enable	Priority
base1	Select	50
base2	Select	50
mgmt2	Select	50

10. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate-5001C interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- fabric3 interface virtual MAC: 00-09-0f-09-00-04
- fabric4 interface virtual MAC: 00-09-0f-09-00-05
- fabric5 interface virtual MAC: 00-09-0f-09-00-06
- mgmt1 keeps its original MAC address
- mgmt2 interface virtual MAC: 00-09-0f-09-00-08
- port1 interface virtual MAC: 00-09-0f-09-00-09
- port2 interface virtual MAC: 00-09-0f-09-00-0a

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP

table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic base1
.
.
.
Current_HWaddr 00:09:0f:09:00:00
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.
```

9. Repeat these steps for the FortiGate-5001C units in chassis slots 4 and 5, with the following differences. Set the mgmt1 interface IP address of each FortiGate-5001C unit to a different IP address.

Set the FortiGate-5001C unit in chassis slot 4 host name to:

New Name	5001C-Slot-4
-----------------	--------------

Set the FortiGate-5001C unit in chassis slot 5 host name to:

New Name	5001C-Slot-5
-----------------	--------------

As you configure each FortiGate, they will negotiate and join the cluster.

To view cluster status

As you add units to the cluster you can log into the GUI of one of the cluster units to view the status of the cluster. The status displays will show each unit as it is added to the cluster.

1. Log into the primary unit or any cluster unit and view the system dashboard.
The System Information dashboard widget shows the **Cluster Name** (example3.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
2. Go to **System > HA** to view the cluster members list.
The list shows three cluster units, their host names, their roles in the cluster, and their priorities. You can check this list at any time to confirm that the cluster is operating normally.

If the cluster members list and the dashboard do not display all of the cluster units, they are not functioning as a cluster.

To troubleshoot the cluster

See [Troubleshooting HA clusters on page 137](#).

To manage each cluster unit

Because you have configured a reserved management interface, you can manage each cluster unit separately by connecting to the IP address you configured for each unit's mgmt1 interface. You can view the status of each cluster unit and make changes to each unit's configuration. For example, as described below, each cluster unit must have its own FortiClient license. You can use the reserved management IP addresses to connect to each cluster unit to install the FortiClient license for that unit.

Usually you would make configuration changes by connecting to the primary unit and changing its configuration. The cluster then synchronizes the configuration changes to all cluster units. If you connect to individual cluster units and change their configuration, those configuration changes are also synchronized to each cluster unit. The exception to this is configuration objects that are not synchronized, such as the host name, FortiClient license and so on.

You can also manage each cluster unit by logging into the primary unit CLI and using the following command to connect to other cluster units:

```
execute ha manage <cluster-index>
```

To add basic configuration settings to the cluster

Use the following steps to configure the cluster.

1. Log into the cluster GUI.
You can log into the primary unit or any one of the cluster units using the appropriate mgmt1 IP address.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Network > Interfaces** and edit the **port1** interface. Set this interface IP address to the address required to connect to the interface to the Internet.
7. Edit the port2 interface and set its IP to an IP address for the internal network.

To add a FortiClient license to each cluster unit

Normally you would add FortiClient licenses to the FortiGates before forming the cluster. However, you can use the following steps to add FortiClient licenses to an operating cluster.

Contact your reseller to purchase FortiClient licenses for your cluster units. Each cluster unit must have its own FortiClient license.

When you receive the license keys you can log into <https://support.fortinet.com> and add a FortiClient license key to each licensed FortiGate. Then, as long as the cluster can connect to the Internet the license keys are downloaded from the FortiGuard network to all of the FortiGates in the cluster.

You can also use the following steps to manually add the license keys to your cluster units from the GUI. Your cluster must be connected to the Internet.

1. Log into the GUI of each cluster unit using its reserved management interface IP address.
2. Go to the **License Information** dashboard widget and beside FortiClient select **Enter License**.
3. Enter the license key and select **OK**.

4. Confirm that the license has been installed and the correct number of FortiClients are licensed.
5. Repeat for all of the cluster units.

You can also use the following command to add the license key from the CLI:

```
execute FortiClient-NAC update-registration-license <license-number>
```

You can connect to the CLIs of each cluster unit using their reserved management IP address.

You can also log into the primary unit CLI and use the `execute ha manage` command to connect to each cluster unit CLI.

Configuring the FortiGate-5000 active-active cluster - CLI

These procedures assume you are starting with three FortiGate-5001C boards and two FortiSwitch-5003B boards installed in a compatible FortiGate-5000 series chassis. The FortiSwitch-5003B boards are in chassis slots 1 and 2 and the FortiGate-5001C boards are in chassis slots 3, 4, and 5 and the chassis is powered on. All devices are in their factory default configuration. No configuration changes to the FortiSwitch-5003B boards are required.

To configure the FortiGate-5005FA2 units

1. From the internal network, log into the CLI of the FortiGate-5001C unit in chassis slot 3 by connecting to the mgmt1 interface.



By default the mgmt1 interface of each FortiGate-5001C unit has the same IP address. To log into each FortiGate-5001C unit separately you could either disconnect the mgmt1 interfaces of the units that you don't want to log into or change the mgmt1 interface IP addresses for each unit by connecting to each unit's CLI from their console port.

You can also use a console connection.

2. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). In this example you will leave **FortiClient** licensing until after you have formed the cluster. Otherwise, all of the FortiGates in a cluster must have the same level of licensing.

We also recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.

3. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. FortiToken licenses can be added at any time because they are synchronized to all cluster members.
4. Change the host name for this FortiGate. For example:

```
config system global
    set hostname 5001C-Slot-3
end
```

5. Enter the following command to display backplane interfaces on the GUI:

```
config system global
    set show-backplane-intf enable
end
```

6. Set the Administrative Status of the base1 and base 2 interfaces to **Up**.

```
config system interface
    edit base1
```



```
        set status up
    next
    edit base2
        set status up
    end
```

7. Add an IP address to the mgmt1 interface.

```
config system interface
    edit mgmt1
        set ip 172.20.120.110/24
        set allowaccess http https ssl ping
    end
```

Because mgmt1 will become the reserved management interface for the cluster unit each FortiGate-5001C should have a different mgmt1 interface IP address. Give the mgmt1 interface an address that is valid for the internal network. Once HA with the reserved Management interface is enabled the IP address of the mgmt1 interface can be on the same subnet as the port2 interface (which will also be connected to the Internal network).

8. Configure HA settings.

```
config system ha
    set mode a-a
    set ha-mgmt-status enable
    set ha-mgmt-interface mgmt1
    set group-name example3.com
    set password HA_pass_3
    set hbdev base1 50 base2 50 mgmt2 50
end
```

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate-5001C interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- fabric3 interface virtual MAC: 00-09-0f-09-00-04
- fabric4 interface virtual MAC: 00-09-0f-09-00-05
- fabric5 interface virtual MAC: 00-09-0f-09-00-06
- mgmt1 keeps its original MAC address
- mgmt2 interface virtual MAC: 00-09-0f-09-00-08
- port1 interface virtual MAC: 00-09-0f-09-00-09
- port2 interface virtual MAC: 00-09-0f-09-00-0a

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```

get hardware nic base1
.
.
.
Current_HWaddr 00:09:0f:09:00:00
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.

```

7. Repeat these steps for the FortiGate-5001C units in chassis slots 4 and 5, with the following differences.

Set the mgmt1 interface IP address of each FortiGate-5001C unit to a different IP address.

Set the FortiGate-5001C unit in chassis slot 4 host name to:

```

config system global
    set hostname 5001C-Slot-4
end

```

Set the FortiGate-5001C unit in chassis slot 5 host name to:

```

config system global
    set hostname 5001C-Slot-5
end

```

As you configure each FortiGate, they will negotiate and join the cluster.

To view cluster status

As you add units to the cluster you can log into the CLI of one of the cluster units using its reserved management interface to view the status of the cluster. The status will show each unit as it is added to the cluster.

For example, the following command output shows the status of the cluster when all three cluster units have been added:

```

get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
Slave : 5001d-slot4      , FG-5KD3914800284
Master: 5001d-slot5      , FG-5KD3914800353
Slave : 5001d-slot3      , FG-5KD3914800344

```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster

See [Troubleshooting HA clusters on page 137](#).

To manage each cluster unit

Because you have configured a reserved management interface, you can manage each cluster unit separately by connecting to the IP address you configured for each unit's mgmt1 interface. You can view the status of each cluster unit and make changes to each unit's configuration. For example, as described below, each cluster unit must have its own FortiClient license. You can use the reserved management IP addresses to connect to each cluster unit to install the FortiClient license for that unit.

Usually you would make configuration changes by connecting to the primary unit and changing its configuration. The cluster then synchronizes the configuration changes to all cluster units. If you connect to individual cluster units and change their configuration, those configuration changes are also synchronized to each cluster unit. The exception to this is configuration objects that are not synchronized, such as the host name, FortiClient license and so on.

You can also manage each cluster unit by logging into the primary unit CLI and using the following command to connect to other cluster units:

```
execute ha manage <cluster-index>
```

To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

To add basic configuration settings to the cluster

Use the following steps to configure the cluster.

1. Log into the cluster CLI.

You can log into the primary unit or any one of the cluster units using the appropriate mgmt1 IP address.

2. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

3. Set the port1 interface IP address to the address required to connect to the interface to the Internet.

```
config system interface
  edit port1
    set ip 10.10.10.10/24
  end
```

4. Set the port2 interface IP address to the address required to connect to the interface to the internal network.

```
config system interface
  edit port2
    set ip 172.20.120.12/24
  end
```

To add a FortiClient license to each cluster unit

Normally you would add FortiClient licenses to the FortiGates before forming the cluster. However, you can use the following steps to add FortiClient licenses to an operating cluster.

Contact your reseller to purchase FortiClient licenses for your cluster units. Each cluster unit must have its own FortiClient license.

When you receive the license keys you can log into <https://support.fortinet.com> and add a FortiClient license key to each licensed FortiGate. Then, as long as the cluster can connect to the Internet the license keys are downloaded from the FortiGuard network to all of the FortiGates in the cluster.

You can also use the following steps to manually add the license keys to your cluster units from the CLI. Your cluster must be connected to the Internet.

1. Log into the CLI of each cluster unit using its reserved management interface IP address.
2. Enter the following command to the unit's serial number:

```
get system status
```

3. Enter the following command to add the license key for that serial number:

```
execute FortiClient-NAC update-registration-license <license-key>
```

4. Confirm that the license has been installed and the correct number of FortiClients are licensed.


```
execute forticlient info
Maximum FortiClient connections: unlimited.
Licensed connections: 114
    NAC: 114
    WANOPT: 0
    Test: 0
Other connections:
    IPsec: 0
    SSLVPN: 0
```

5. Repeat for all of the cluster units.

You can also log into the primary unit CLI and use the `execute ha manage` command to connect to each cluster unit CLI.

Converting a standalone FortiGate to a cluster

In this recipe, a backup FortiGate will be installed and connected to a FortiGate that has previously been installed to provide redundancy if the primary FortiGate fails.



1. Adding the backup FortiGate and configuring HA


If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license.

If you have not already done so, register the primary FortiGate and apply licenses to it before setting up the cluster. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMs). All of the FortiGates in a cluster must have the same level of licensing. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.

License Information

Support Contract	Registration	Registered (bdictie@fortinet.com)	Launch Portal
FortiGuard	IPS & Application Control	Licensed (Expires 2016-08-22)	
	AntiVirus	Licensed (Expires 2016-08-22)	
	Web Filtering	Licensed (Expires 2016-08-21)	
	Anti-Spam Filtering	Licensed (Expires 2016-08-21)	
FortiCloud	Account		Activate
FortiSandbox	FortiSandbox Appliance	Not Configured	Configure
FortiClient	Status	Free License	How to Purchase
	Clients Registered	0 of 10	Enter License
	FortiClient Installers		Details
FortiToken Mobile	Tokens Assigned	0 of 2	

Connect your network as shown in the initial diagram, with Ethernet cables connecting the **HA** heartbeat interfaces of the two FortiGates. If your FortiGate does not have dedicated HA heartbeat interfaces, you can use different interfaces, provided they are not used for any other function.



A switch must be used between the FortiGates and Internet, and another is required between the FortiGates and the internal network, as shown in the network diagram for this recipe.

Connect to the primary FortiGate and go to **System > Dashboard > Status** and locate the **System Information** widget.

Current Name FG100D3G12804410
New Name

Change the unit's **Host Name** to identify it as the primary FortiGate.

In the **System Information** widget, configure **HA Status**. Set the **Mode** to **Active-Passive** and set a **Group Name** and **Password**.

Ensure that the two **Heartbeat Interfaces** are selected and their priorities are both set to 50.

Mode

Active-Passive

Device Priority

128

☐ Reserve Management Port for Cluster Member

Internal

Cluster Settings

Group Name

HA-cluster

Password

.....

☐ Enable Session Pick-up

	Port Monitor	Heartbeat Interface	
		Enable	Priority(0-512)
dmz	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
ha1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<div>50</div>
ha2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<div>50</div>
mgmt	<input type="checkbox"/>		
port9	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
port10	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
port11	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
port14	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
port15	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
port16	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
wan1	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>
wan2	<input type="checkbox"/>	<input type="checkbox"/>	<div>0</div>

Connect to the backup FortiGate and go to **System > Dashboard > Status**.

Change the unit's **Host Name** to identify it as the backup FortiGate.

Current Name

FG100D3G12801361

New Name

Backup_FortiGate

Configure **HA Status** and set the **Mode** to **Active-Passive**.

Set the **Device Priority** to be lower than the primary FortiGate. Ensure that the **Group Name** and **Password** match those on the primary FortiGate.

Ensure that the two **Heartbeat Interfaces** are selected and their priorities are both set to 50.

Mode Active-Passive

Device Priority 50

☐ Reserve Management Port for Cluster Member dmz

Cluster Settings

Group Name HA-cluster

Password *****

☐ Enable Session Pick-up

	Port Monitor	Heartbeat Interface	
		Enable	Priority(0-512)
dmz	<input type="checkbox"/>	<input type="checkbox"/>	0
ha1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
ha2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	50
mgmt	<input type="checkbox"/>		
port1	<input type="checkbox"/>	<input type="checkbox"/>	0
wan1	<input type="checkbox"/>	<input type="checkbox"/>	0
wan2	<input type="checkbox"/>	<input type="checkbox"/>	0

Connect to the primary FortiGate and go to **System > HA** to view the cluster information.

	Cluster Member	Hostname	Serial No.	Role	Priority
		Primary_FortiGate	FG100D3G12804410	MASTER	128
		Backup_FortiGate	FG100D3G12801361	SLAVE	50

Select **View HA Statistics** for more information on how the cluster is operating and processing traffic.

Unit	Status	Up Time	Monitor		
Primary_FortiGate FG100D3G12804410		0 days		Active Sessions	Virus Detected
		1 hours		26	0
		44 minutes		Total Packets	Intrusion Detected
		2 seconds		81857	0
Backup_FortiGate FG100D3G12801361		2 days		Network Utilization	Total Bytes
		0 hours		78 Kbps	27300058
		15 minutes		Active Sessions	Virus Detected
		15 seconds		6	0
				Total Packets	Intrusion Detected
				13 Kbps	2778691497
					0

2. Results

Normally, traffic should now be flowing through the primary FortiGate. However, if the primary FortiGate is unavailable, traffic should failover and the backup FortiGate will be used. Failover will also cause the primary and backup FortiGates to reverse roles, even when both FortiGates are available again.

To test this, ping the IP address 8.8.8.8 using a PC on the internal network. After a moment, power off the primary FortiGate. You will see a momentary pause in the Ping results, until traffic diverts to the backup FortiGate, allowing the Ping traffic to continue.

```
Reply from 8.8.8.8: bytes=32 time=50ms TTL=53
Reply from 8.8.8.8: bytes=32 time=38ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Request timed out.
Reply from 8.8.8.8: bytes=32 time=482ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Reply from 8.8.8.8: bytes=32 time=36ms TTL=53
Reply from 8.8.8.8: bytes=32 time=37ms TTL=53
Reply from 8.8.8.8: bytes=32 time=38ms TTL=53
```

If you are using port monitoring, you can also unplug the primary FortiGate's Internet-facing interface to test failover.

Replacing a failed cluster unit

This procedure describes how to remove a failed cluster unit from a cluster and add a new one to replace it. You can also use this procedure to remove a failed unit from a cluster, repair it and add it back to the cluster. Replacing a failed does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the replacement unit.

You can use this procedure to replace more than one cluster unit.

To replace a failed cluster unit

1. Disconnect the failed unit from the cluster and the network.
If you maintain other connections between the network and the still functioning cluster unit or units and between remaining cluster units network traffic will continue to be processed.
2. Repair the failed cluster unit, or obtain a replacement unit with the exact same hardware configuration as the failed cluster unit.
3. Install the same firmware build on the repaired or replacement unit as is running on the cluster.
4. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information			
Support Contract	Registration	✓ Registered (bdickie@fortinet.com)	Launch Portal
FortiGuard	IPS & Application Control	✓ Licensed (Expires 2016-08-22)	
	AntiVirus	✓ Licensed (Expires 2016-08-22)	
	Web Filtering	✓ Licensed (Expires 2016-08-21)	
	Anti-Spam Filtering	✓ Licensed (Expires 2016-08-21)	
FortiCloud	Account		Activate
FortiSandbox	FortiSandbox Appliance	✗ Not Configured	Configure
FortiClient	Status	✓ Free License	How to Purchase
	Clients Registered	0 of 10	Enter License
	FortiClient Installers		Details
FortiToken Mobile	Tokens Assigned	0 of 2	

5. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
6. Configure the repaired or replacement unit for HA operation with the same HA configuration as the cluster.
7. If the cluster is running in Transparent mode, change the operating mode of the repaired or replacement unit to Transparent mode.
8. Connect the repaired or replacement cluster unit to the cluster.
For an example see [How to set up FGCP clustering \(recommended steps\) on page 66](#).
9. Power on the repaired or replacement cluster unit.
When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the repaired or replacement unit configuration with the configuration of the primary unit.

You can add a repaired or replacement unit to a functioning cluster at any time. The repaired or replacement cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.

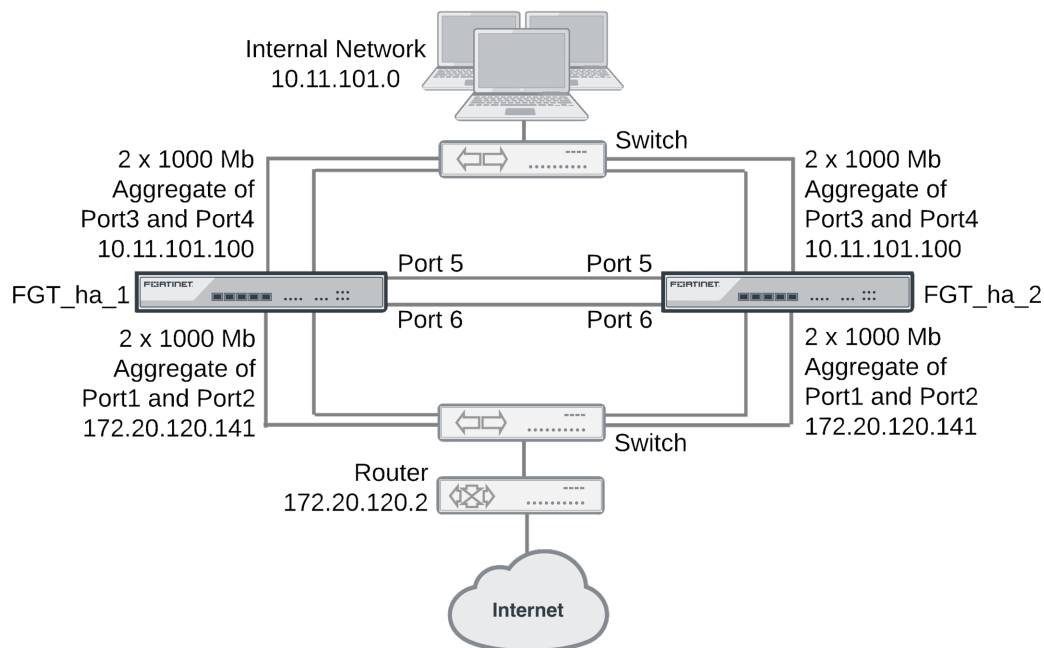
FGCP HA with 802.3ad aggregated interfaces

On FortiGate models that support it you can use 802.3ad link aggregation to combine two or more interfaces into a single aggregated interface. 802.3ad Link Aggregation and its management protocol, Link Aggregation Control Protocol (LACP) are a method for combining multiple physical links into a single logical link. This increases both potential throughput and network resiliency. Using LACP, traffic is distributed among the physical interfaces in the link, potentially resulting in increased performance.

This example describes how to configure an HA cluster consisting of two FortiGates with two aggregated 1000 Mb connections to the Internet using port1 and port2 and two aggregated 1000 Mb connections to the internal network using port3 and port4. The aggregated interfaces are also configured as HA monitored interfaces.

Each of the aggregate links connects to a different switch. Each switch is configured for link aggregation (2x1000Mb).

Example cluster with aggregate interfaces



HA interface monitoring, link failover, and 802.3ad aggregation

When monitoring the aggregated interface, HA interface monitoring treats the aggregated link as a single interface and does not monitor the individual physical interfaces in the link. HA interface monitoring registers the link to have failed only if all the physical interfaces in the link have failed. If only some of the physical interfaces in the link fail or become disconnected, HA considers the link to be operating normally.

HA MAC addresses and 802.3ad aggregation

If a configuration uses the Link Aggregate Control Protocol (LACP) (either passive or active), LACP is negotiated over all of the interfaces in any link. For a standalone FortiGate, the FortiGate LACP implementation uses the MAC address of the first interface in the link to uniquely identify that link. For example, a link consisting of port1 and port2 interfaces would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. An aggregate interface in a cluster acquires the virtual MAC address that would have been acquired by the first interface in the aggregate.

Link aggregation, HA failover performance, and HA mode

To operate an active-active or active-passive cluster with aggregated interfaces and for best performance of a cluster with aggregated interfaces, the switches used to connect the cluster unit aggregated interfaces together should support configuring multiple Link Aggregation (LAG) groups.

For example, the cluster shown above should be configured into two LAG groups on the external switch: one for the port1 and port2 aggregated interface of FGT_ha_1 and a second one for the port1 and port2 aggregate interface of FGT_ha_2. You should also be able to do the same on the internal switch for the port3 and port4 aggregated interfaces of each cluster unit.

As a result, the subordinate unit aggregated interfaces would participate in LACP negotiation while the cluster is operating. In an active-active mode cluster, packets could be redirected to the subordinate unit interfaces. As well, in active-active or active-passive mode, after a failover the subordinate unit can become a primary unit without having to perform LACP negotiation before it can process traffic. Performing LACP negotiation causes a minor failover delay.

However if you cannot configure multiple LAG groups on the switches, due to the primary and subordinate unit interfaces having the same MAC address, the switch will put all of the interfaces into the same LAG group which would disrupt the functioning of the cluster. To prevent this from happening, you must change the FortiGate aggregated interface configuration to prevent subordinate units from participating in LACP negotiation.

For example, use the following command to prevent subordinate units from participating in LACP negotiation with an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```

As a result of this setting, subordinate unit aggregated interfaces cannot accept packets. This means that you cannot operate the cluster in active-active mode because in active-active mode the subordinate units must be able to receive and process packets. Also, failover may take longer because after a failover the subordinate unit has to perform LACP negotiation before being able to process network traffic.

Also, it may also be necessary to configure the switch to use Passive or even Static mode for LACP to prevent the switch from sending packets to the subordinate unit interfaces, which won't be able to process them.

Finally, in some cases depending on the LACP configuration of the switches, you may experience delayed failover if the FortiGate LACP configuration is not compatible with the switch LACP configuration. For example, in some cases setting the FortiGate LACP mode to static reduces the failover delay because the FortiGate does not perform LACP negotiation. However there is a potential problem with this configuration because static LACP does not send periodic LAC Protocol Data Unit (LACPDU) packets to test the connections. So a non-physical

failure (for example, if a device is not responding because its too busy) may not be detected and packets could be lost or delayed.

General configuration steps

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. View cluster status.
4. Add basic configuration settings and configure the aggregated interfaces.
 - Add a password for the admin administrative account.
 - Add the aggregated interfaces.
 - Disable `lacp-ha-slave` so that the subordinate unit does not send LACP packets.
 - Add a default route.

You could also configure aggregated interfaces in each FortiGate before the units form a cluster.

5. Configure HA port monitoring for the aggregated interfaces.

Configuring active-passive HA cluster that includes aggregated interfaces - GUI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information

	Support Contract	Registration	✓ Registered (bdickie@fortinet.com)	Launch Portal
	FortiGuard	IPS & Application Control	✓ Licensed (Expires 2016-08-22)	
		AntiVirus	✓ Licensed (Expires 2016-08-22)	
		Web Filtering	✓ Licensed (Expires 2016-08-21)	
		Anti-Spam Filtering	✓ Licensed (Expires 2016-08-21)	
	FortiCloud	Account		Activate
	FortiSandbox	FortiSandbox Appliance	✗ Not Configured	Configure
	FortiClient	Status	✓ Free License	How to Purchase
		Clients Registered	<input type="text" value="0 of 10"/>	Enter License
		FortiClient Installers	<input type="text" value="0 of 2"/>	Details
	FortiToken Mobile	Tokens Assigned	<input type="text" value="0 of 2"/>	

- You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
- On the **System Information** dashboard widget, beside **Host Name** select **Change**.
- Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

- Select **OK**.
- Go to **System > HA** and change the following settings.

Mode	Active-Passive	
Group Name	example5.com	
Password	HA_pass_5	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration to not use those interfaces.

- Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate.
8. Repeat these steps for the second FortiGate.

Set the second FortiGate host name to:

New Name	FGT_ha_2
-----------------	----------

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up an aggregated interface and port1 and port2 of FGT_ha_2 make up a second aggregated interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up an aggregated interface and port3 and port4 of FGT_ha_2 make up another aggregated interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
The units negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete, the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.
The System Information dashboard widget shows the **Cluster Name** (example5.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
2. Go to **System > HA** to view the cluster members list.
The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Admin > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Router > Static > Static Routes** and temporarily delete the default route.

You cannot add an interface to a aggregated interface if any settings (such as the default route) are configured for it.

7. Go to **System > Network > Interfaces** and select **Create New** to add the aggregate interface to connect to the Internet.
8. Set **Type** to **802.3ad Aggregate** and configure the aggregate interface to be connected to the Internet:

Name	Port1_Port2
Physical Interface Members	port1, port2
IP/Network Mask	172.20.120.141/24

9. Select **OK**.
10. Select **Create New** to add the aggregate interface to connect to the internal network.
11. Set **Type** to **802.3ad Aggregate** and configure the aggregate interface to be connected to the Internet:

Name	Port3_Port4
Physical Interface Members	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12. Select **OK**.

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11

- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Connect to the CLI and enter the following command to disable sending LACP packets from the subordinate unit:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  next
  edit Port3_Port4
    set lacp-ha-slave disable
end
```

14. Go to **Router > Static > Static Routes**.

15. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

16. Select **OK**.

To configure HA port monitoring for the aggregate interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the aggregate interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4. Select **OK**.

Configuring active-passive HA cluster that includes aggregate interfaces - CLI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. Install any third-party certificates on the FortiGate.
3. Change the host name for this FortiGate:

```
config system global
  set hostname FGT_ha_1
```

```
end
```

4. Configure HA settings.

```
config system ha
    set mode a-p
    set group-name example5.com
    set password HA_pass_5
    set hbdev port5 50 port6 50
end
```

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
```

```
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Repeat these steps for the other FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up an aggregated interface and port1 and port2 of FGT_ha_2 make up another aggregated interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up an interfaced and port3 and port4 of FGT_ha_2 make up another aggregated interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Enter the following command to view the status of the cluster:

```
get system ha status
HA Health Status: OK
```

```
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 137](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings and the aggregate interfaces.

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

2. Temporarily delete the default route.

You cannot add an interface to an aggregate interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
  delete 1
end
```

3. Add the aggregate interfaces:

```
config system interface
  edit Port1_Port2
    set type aggregate
    set lacp-ha-slave disable
    set member port1 port2
    set ip 172.20.120.141/24
    set vdom root
  next
  edit Port3_Port4
    set type aggregate
    set lacp-ha-slave disable
    set member port3 port4
    set ip 10.11.101.100/24
    set vdom root
  end
```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and

port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4. Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the aggregate interfaces


1. Configure HA port monitoring for the aggregate interfaces.

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Example HA and redundant interfaces

On FortiGate models that support it you can combine two or more interfaces into a single redundant interface. A redundant interface consists of two or more physical interfaces. Traffic is processed by the first physical interface in the redundant interface. If that physical interface fails, traffic fails over to the next physical interface. Redundant interfaces don't have the benefit of improved performance that aggregate interfaces can have, but they do provide failover if a physical interface fails or is disconnected.

Example cluster with a redundant interfaces



This example describes how to configure an HA cluster consisting of two FortiGates with a a redundant interface connection to the Internet and to an internal network. The connection to the Internet uses port1 and port2. The connection to the internal network uses port3 and port4. The HA heartbeat uses port5 and port6.

The redundant interfaces are also configured as HA monitored interfaces.

HA interface monitoring, link failover, and redundant interfaces

HA interface monitoring monitors the redundant interface as a single interface and does not monitor the individual physical interfaces in the redundant interface. HA interface monitoring registers the redundant interface to have failed only if all the physical interfaces in the redundant interface have failed. If only some of the physical interfaces in the redundant interface fail or become disconnected, HA considers the redundant interface to be operating normally.

HA MAC addresses and redundant interfaces

For a standalone FortiGate a redundant interface has the MAC address of the first physical interface added to the redundant interface configuration. A redundant interface consisting of port1 and port2 would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. A redundant interface in a cluster acquires the virtual MAC address that would have been acquired by the first physical interface added to the redundant interface configuration.

Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode

HA assigns the same virtual MAC addresses to the subordinate unit interfaces as are assigned to the corresponding primary unit interfaces. Consider a cluster of two FortiGates operating in active-passive mode with

a redundant interface consisting of port1 and port2. You can connect multiple redundant interfaces to the same switch if you configure the switch so that it defines multiple separate redundant interfaces and puts the redundant interfaces of each cluster unit into separate redundant interfaces. In this configuration, each cluster unit forms a separate redundant interface with the switch.

However, if the switch is configured with a single four-port redundant interface configuration, because the same MAC addresses are being used by both cluster units, the switch adds all four interfaces (port1 and port2 from the primary unit and port1 and port2 from the subordinate unit) to the same redundant interface.

To avoid unpredictable results, when you connect a switch to multiple redundant interfaces in an active-passive cluster you should configure separate redundant interfaces on the switch; one for each cluster unit.

Connecting multiple redundant interfaces to one switch while operating in active-active HA mode

In an active-active cluster, all cluster units send and receive packets. To operate a cluster with redundant interfaces in active-active mode, with multiple redundant interfaces connected to the same switch, you must separate the redundant interfaces of each cluster unit into different redundant interfaces on the connecting switch.

General configuration steps

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. View cluster status.
4. Add basic configuration settings and configure the redundant interfaces.
 - Add a password for the admin administrative account.
 - Add the redundant interfaces.
 - Add a default route.

You could also configure redundant interfaces in each FortiGate before they form a cluster.

5. Configure HA port monitoring for the redundant interfaces.

Configuring active-passive HA cluster that includes redundant interfaces - GUI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a

cluster must have the same level of licensing.

Category	Item	Status	Action
Support Contract	Registration	Registered (bdickie@fortinet.com)	Launch Portal
FortiGuard	IPS & Application Control	Licensed (Expires 2016-08-22)	
	AntiVirus	Licensed (Expires 2016-08-22)	
	Web Filtering	Licensed (Expires 2016-08-21)	
	Anti-Spam Filtering	Licensed (Expires 2016-08-21)	
FortiCloud	Account		Activate
FortiSandbox	FortiSandbox Appliance	Not Configured	Configure
FortiClient	Status	Free License	How to Purchase
	Clients Registered	0 of 10	Enter License
	FortiClient Installers		Details
FortiToken Mobile	Tokens Assigned	0 of 2	

- You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
- On the **System Information** dashboard widget, beside **Host Name** select **Change**.
- Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

- Select **OK**.
- Go to **System > HA** and change the following settings.

Mode	Active-Passive	
Group Name	example6.com	
Password	HA_pass_6	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat

configuration.

7. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate.

8. Repeat these steps for the second FortiGate.

Set the second FortiGate host name to:

New Name	FGT_ha_2
-----------------	----------

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up a redundant interface and port1 and port2 of FGT_ha_2 make up another redundant interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up a redundant interface and port3 and port4 of FGT_ha_2 make up another redundant interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.
The System Information dashboard widget shows the **Cluster Name** (example5.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
2. Go to **System > HA** to view the cluster members list.
The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 137](#).

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Admin > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.

5. Select **OK**.
6. Go to **Router > Static > Static Routes** and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to **System > Network > Interfaces** and select **Create New** to add the redundant interface to connect to the Internet.
8. Set **Type** to **Redundant Interface** and configure the redundant interface to be connected to the Internet:

Name	Port1_Port2
Physical Interface Members	port1, port2
IP/Netmask	172.20.120.141/24

9. Select **OK**.
10. Select **Create New** to add the redundant interface to connect to the internal network.
11. Set **Type** to **Redundant Interface** and configure the redundant interface to be connected to the Internet:

Name	Port3_Port4
Physical Interface Members	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12. Select **OK**.
The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f

- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Go to **Router > Static > Static Routes**.

14. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

15. Select **OK**.

To configure HA port monitoring for the redundant interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the redundant interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4. Select **OK**.

Configuring active-passive HA cluster that includes redundant interfaces - CLI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. Change the host name for this FortiGate:

```
config system global
    set hostname FGT_ha_1
end
```

4. Configure HA settings.

```
config system ha
```

```

set mode a-p
set group-name example6.com
set password HA_pass_6
set hbdev port5 50 port6 50
end

```

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```

get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.

```

- Repeat these steps for the other FortiGate.

Set the other FortiGate host name to:

```
config system global
  set hostname FGT_ha_2
end
```

To connect the cluster to the network

- Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up a redundant interface and port1 and port2 of FGT_ha_2 make up another redundant interface.
- Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up a redundant interface and port3 and port4 of FGT_ha_2 make up another redundant interface.
- Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.
When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

- Log into the CLI.
- Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

- Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
```

```

Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.

```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 137](#).

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings and the redundant interfaces.

1. Add a password for the admin administrative account.

```

config system admin
  edit admin
    set password <psswr>
  end

```

2. Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```

config router static
  delete 1
end

```

3. Add the redundant interfaces:

```

config system interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
    set ip 172.20.120.141/24
    set vdom root
  next
  edit Port3_Port4
    set type redundant
    set member port3 port4
    set ip 10.11.101.100/24
    set vdom root
  end

```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4. Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the redundant interfaces

1. Configure HA port monitoring for the redundant interfaces.

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Troubleshooting HA clusters

This section describes some HA clustering troubleshooting techniques.

Ignoring hardware revisions

Many FortiGate platforms have gone through multiple hardware versions and in some cases the hardware changes prevent cluster formation. If you run into this problem you can use the following command on each FortiGate to cause the cluster to ignore different hardware versions:

```
execute ha ignore-hardware-revision enable
```

This command is only available on FortiGates that have had multiple hardware revisions.

By default the command is set to prevent cluster formation between FortiGates with different hardware revisions. You can enter the following command to view its status:

```
execute ha ignore-hardware-revision status
```

Usually the incompatibility is caused by different hardware versions having different hard disks and enabling this command disables each FortiGate's hard disks. As a result of disabling hard disks the cluster will not support logging to the hard disk or WAN Optimization.

If the FortiGates do have compatible hardware versions or if you want to run a FortiGate in standalone mode you can enter the following command to disable ignoring the hardware revision and enable the hard disks:

```
execute ha ignore-hardware-revision disable
```

Affected models include but are not limited to:

- FortiGate-100D
- FortiGate-300C
- FortiGate-600C
- FortiGate-800C
- FortiGate-80C and FortiWiFi-80C
- FortiGate-60C



It's possible that a cluster will not form because the disk partition sizes of the cluster units are different. You can use the `diagnose sys ha checksum test | grep storage` command to check the disk storage checksum of each cluster unit. If the checksums are different then contact Fortinet support for help in setting up compatible storage partitions.

Before you set up a cluster

Before you set up a cluster ask yourself the following questions about the FortiGates that you are planning to use to create a cluster.

1. Do all the FortiGates have the same hardware configuration? Including the same hard disk configuration?
2. Do all of the FortiGates have the same FortiGuard, FortiCloud, FortiClient, VDOM and FortiOS Carrier licensing?
3. Do all the FortiGates have the same firmware build?
4. Are all the FortiGates set to the same operating mode (NAT or Transparent)?
5. Are all the FortiGates operating in single VDOM mode?
6. If the FortiGates are operating in multiple VDOM mode do they all have the same VDOM configuration?



In some cases you may be able to form a cluster if different FortiGates have different firmware builds, different VDOM configurations, and are in different operating modes. However, if you encounter problems they may be resolved by installing the same firmware build on each unit, and give them the same VDOM configuration and operating mode. If the FortiGates in the cluster have different licenses, the cluster will form but it will operate with the lowest licensing level.

Troubleshooting the initial cluster configuration

This section describes how to check a cluster when it first starts up to make sure that it is configured and operating correctly. This section assumes you have already configured your HA cluster.

To verify that a cluster can process traffic and react to a failure

1. Add a basic security policy configuration and send network traffic through the cluster to confirm connectivity.
For example, if the cluster is installed between the Internet and an internal network, set up a basic internal to external security policy that accepts all traffic. Then from a PC on the internal network, browse to a website on the Internet or ping a server on the Internet to confirm connectivity.
2. From your management PC, set ping to continuously ping the cluster, and then start a large download, or in some other way establish ongoing traffic through the cluster.
3. While traffic is going through the cluster, disconnect the power from one of the cluster units.
You could also shut down or restart a cluster unit.
Traffic should continue with minimal interruption.
4. Start up the cluster unit that you disconnected.
The unit should re-join the cluster with little or no affect on traffic.
5. Disconnect a cable from one of the HA heartbeat interfaces.
The cluster should keep functioning, using the other HA heartbeat interface.
6. If you have port monitoring enabled, disconnect a network cable from a monitored interface.
Traffic should continue with minimal interruption.

To verify the cluster configuration from the GUI

Use these steps if a cluster is formed just to verify its status and configuration.

1. Log into the cluster GUI.
2. Check the system dashboard to verify that the System Information widget displays all of the cluster units.
3. Check the Unit Operation widget graphic to verify that the correct cluster unit interfaces are connected.
4. Go to **System > HA** or from the System Information dashboard widget select **HA Status > Configure** and verify that all of the cluster units are displayed on the HA Cluster list.
5. From the cluster members list, edit the primary unit (master) and verify the cluster configuration is as expected.

To troubleshoot the cluster configuration from the GUI

Use these steps if the FortiGates don't successfully form a cluster:

1. Connect to each cluster unit GUI and verify that the HA configurations are the same. The HA configurations of all of the cluster units must be identical. Even though the HA configuration is very simple you can easily make a small mistake that prevents a FortiGate from joining a cluster.
2. If the configurations are the same, try re-entering the HA **Password** on each cluster unit in case you made an error typing the password when configuring one of the cluster units.
3. Check that the correct interfaces of each cluster unit are connected.
Check the cables and interface LEDs.

Use the Unit Operation dashboard widget, system network interface list, or cluster members list to verify that each interface that should be connected actually is connected.

If the link is down re-verify the physical connection. Try replacing network cables or switches as required.

To verify the cluster configuration from the CLI

Use these steps if a cluster is formed just to verify its status and configuration.

1. Log into each cluster unit CLI.
You can use the console connection if you need to avoid the problem of units having the same IP address.
2. Enter the command `get system status`.
Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Verify that the `get system ha status` command shows that the cluster health is OK and shows that all of the cluster units have joined the cluster.
4. Enter the `get system ha` command to verify that the HA configuration is correct and the same for each cluster unit.

To troubleshoot the cluster configuration from the CLI

Try these steps if the FortiGates don't successfully form a cluster:

1. Try using the following command to re-enter the cluster password on each cluster unit in case you made an error typing the password when configuring one of the cluster units.

```
config system ha
    set password <password>
end
```

2. Check that the correct interfaces of each cluster unit are connected.
Check the cables and interface LEDs.

Use `get hardware nic <interface_name>` command to confirm that each interface is connected. If the interface is connected the command output should contain a `Link: up` entry similar to the following:

```
get hardware nic port1
.
.
.
Link: up
.
.
```

If the link is down, re-verify the physical connection. Try replacing network cables or switches as required.

More troubleshooting information

Much of the information in this HA guide can be useful for troubleshooting HA clusters. Here are some links to sections with more information.

- If sessions are lost after a failover you may need to change route-ttl to keep synchronized routes active longer. See [Synchronizing kernel routing tables on page 239](#)
- To control which cluster unit becomes the primary unit, you can change the device priority and enable override. See [Controlling primary unit selection using device priority and override on page 49](#)
- Changes made to a cluster can be lost if override is enabled. See [Configuration changes can be lost if override is enabled on page 49](#)
- When override is enabled, after a failover traffic may be disrupted if the primary unit rejoins the cluster before the session tables are synchronized or for other reasons such as if the primary unit is configured for DHCP or PPPoE. See [Delaying how quickly the primary unit rejoins the cluster when override is enabled on page 50](#).
- In some cases, age differences among cluster units result in the wrong cluster unit becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units. You can resolve this problem by resetting the age of one or more cluster units. See [Primary unit selection and age on page 41](#) You can also adjust how sensitive the cluster is to age differences. This can be useful if large age differences cause problems. See [Cluster age difference margin \(grace period\) on page 42](#) and [Changing the cluster age difference margin on page 42](#).
- If one of the cluster units needs to be serviced or removed from the cluster for other reasons, you can do so without affecting the operation of the cluster. See [Disconnecting a cluster unit from a cluster on page 211](#)
- The GUI and CLI will not allow you to configure HA if you have enabled FGSP HA. See [FortiGate Session Life Support Protocol \(FGSP\) on page 298](#).
- The GUI and CLI will not allow you to configure HA if one or more FortiGate interfaces is configured as a PPTP or L2TP client.
- The FGCP is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. See [FortiGate HA compatibility with DHCP and PPPoE on page 50](#).
- Some third-party network equipment may prevent HA heartbeat communication, resulting in a failure of the cluster or the creation of a split brain scenario. For example, some switches use packets with the same Ethertype as HA heartbeat packets use for internal functions and when used for HA heartbeat communication the switch generates CRC errors and the packets are not forwarded. See [Heartbeat packet Ethertypes on page 221](#).
- Very busy clusters may not be able to send HA heartbeat packets quickly enough, also resulting in a split brain scenario. You may be able to resolve this problem by modifying HA heartbeat timing. See [Modifying heartbeat timing on page 222](#).
- Very busy clusters may suffer performance reductions if session pickup is enabled. If possible you can disable this feature to improve performance. If you require session pickup for your cluster, several options are available for improving session pickup performance. See [Session failover \(session-pickup\) on page 262](#).
- If it takes longer than expected for a cluster to failover you can try changing how the primary unit sends gratuitous ARP packets. See [Changing how the primary unit sends gratuitous ARP packets after a failover on page 225](#).

- You can also improve failover times by configuring the cluster for subsecond failover. See [Subsecond failover on page 249](#) and [Failover performance on page 260](#).
- When you first put a FortiGate in HA mode you may lose connectivity to the unit. This occurs because HA changes the MAC addresses of all FortiGate interfaces, including the one that you are connecting to. The cluster MAC addresses also change if you change some HA settings such as the cluster group ID. The connection will be restored in a short time as your network and PC updates to the new MAC address. To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all ARP table entries). You may be able to delete the ARP table of your management PC from a command prompt using a command similar to `arp -d`.
- Since HA changes all cluster unit MAC addresses, if your network uses MAC address filtering you may have to make configuration changes to account for the HA MAC addresses.
- A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID. See [Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain on page 229](#).
- The cluster CLI displays `slave is not in sync` messages if there is a synchronization problem between the primary unit and one or more subordinate units. See [How to diagnose HA out of sync messages on page 238](#).
- If you have configured dynamic routing and the new primary unit takes too long to update its routing table after a failover you can configure graceful restart and also optimize how routing updates are synchronized. See [Configuring graceful restart for dynamic routing failover on page 241](#) and [Synchronizing kernel routing tables on page 239](#).
- Some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur after a link failover if the switch does not detect the failure and does not clear its MAC forwarding table. See [Updating MAC forwarding tables when a link failover occurs on page 247](#).
- If a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails you can enable remote link failover to maintain communication. See [Remote link failover on page 250](#).
- If you find that some cluster units are not running the same firmware build you can reinstall the correct firmware build on the cluster to upgrade all cluster units to the same firmware build. See [Synchronizing the firmware build running on a new cluster unit on page 197](#).

Virtual clusters

This chapter provides an introduction to virtual clustering and also contains general procedures and configuration examples that describe how to configure FortiGate HA virtual clustering.

Virtual clustering overview

Virtual clustering is an extension of the FGCP for a cluster of two FortiGates operating with multiple VDOMs enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

The figure below shows an example virtual cluster configuration consisting of two FortiGates. The virtual cluster has two virtual domains, root and Eng_vdm.

The root virtual domain includes the port1 and port2 interfaces. The Eng_vdm virtual domain includes the port5 and port6 interfaces. The port3 and port4 interfaces (not shown in the diagram) are the HA heartbeat interfaces.



This chapter describes setting up FortiGate virtual clustering with a cluster of two FortiGates with multiple VDOMs enabled. You can also create a cluster of three or four FortiGates operating with multiple VDOMs. A future version of this guide will include a configuration example of this setup.

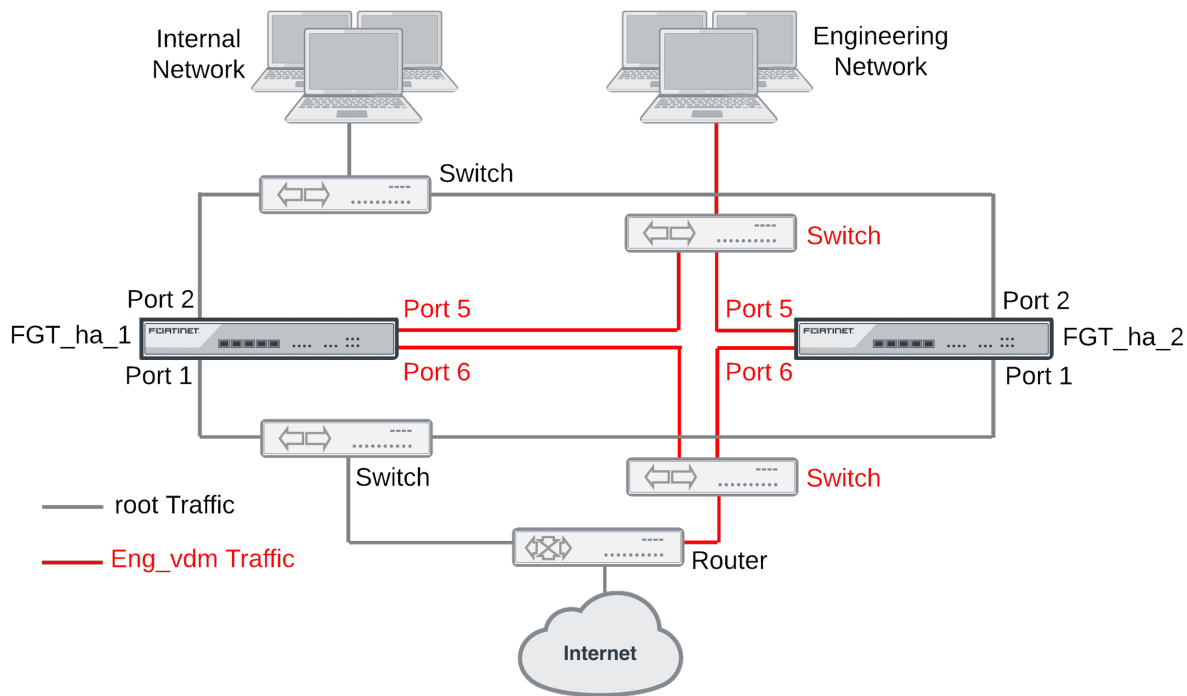
Virtual clustering and failover protection

Virtual clustering operates on a cluster of two FortiGates with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGates in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

Example virtual cluster



Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM partitioning from the GUI by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI `config system ha` command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.



If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.

For more information about HA override see [HA override on page 47](#).

Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all

sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved.

From the GUI you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.



The device priorities for virtual cluster 1 and virtual cluster 2 are not synchronized between the FortiGates in the virtual cluster. You must configure these device priorities separately for each cluster unit.

From the CLI you configure VDOM partitioning by setting the HA mode to `a-p`. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1 (`vcluster 1`). You do the same for virtual cluster 2 (`vcluster 2`) by entering the `config secondary-vcluster` command.



If your cluster has a VLAN that is part of a different VDOM than the physical interface that the VLAN has been added to, then you must configure VDOM partitioning to keep both of these VDOMs in the same virtual cluster (`vcluster`).

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit.

If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

Configuring HA for virtual clustering

If your cluster uses VDOMs, you are configuring virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below.

To configure HA options for a cluster with VDOMs enabled:

- Log into the global GUI and go to **System > HA**.
- From the CLI, log into the Global Configuration:

The following example shows how to configure active-active virtual clustering:

```
config global
  config system ha
    set mode a-a
    set group-name vexample1.com
    set password vHA_pass_1
  end
end
```

The following example shows how to configure active-passive virtual clustering:

```
config global
  config system ha
    set mode a-p
    set group-name vexample1.com
    set password vHA_pass_1
  end
end
```

The following example shows how to configure VDOM partitioning for virtual clustering. In the example, the FortiGate is configured with three VDOMs (domain_1, domain_2, and domain_3) in addition to the root VDOM. The example shows how to set up a basic HA configuration that sets the device priority of virtual cluster 1 to 200. The example also shows how to enable `vcluster2`, how to set the device priority of virtual cluster 2 to 100 and how to add the virtual domains `domain_2` and `domain_3` to virtual cluster 2.

When you enable multiple VDOMs, `vcluster2` is enabled by default. Even so the command to enable `vcluster2` is included in this example in case for some reason it has been disabled. When `vcluster2` is enabled, `override` is also enabled.

The result of this configuration would be that the cluster unit that you are logged into becomes the primary unit for virtual cluster 1. This cluster unit processes all traffic for the root and domain_1 virtual domains.

```
config global
  config system ha
    set mode a-p
    set group-name vexample1.com
    set password vHA_pass_1
    set priority 200
    set vcluster2 enable
    config secondary-vcluster
      set vdom domain_2 domain_3
      set priority 100
    end
  end
end
```

The following example shows how to use the `execute ha manage` command to change the device priorities for virtual cluster 1 and virtual cluster 2 for the other unit in the cluster. The commands set the device priority of virtual cluster 1 to 100 and virtual cluster 2 to 200.

The result of this configuration would be that the other cluster unit becomes the primary unit for virtual cluster 2. This other cluster unit would process all traffic for the domain_2 and domain_3 virtual domains.

```
config global
  execute ha manage 1
  config system ha
    set priority 100
```

```
        set vcluster2 enable
        config secondary-vcluster
            set priority 200
        end
    end
end
end
end
```

Example virtual clustering with two VDOMs and VDOM partitioning

This section describes how to configure the example virtual clustering configuration shown below. This configuration includes two virtual domains, root and Eng_vdm and includes VDOM partitioning that sends all root VDOM traffic to FGT_ha_1 and all Eng_vdm VDOM traffic to FGT_ha_2. The traffic from the internal network and the engineering network is distributed between the two FortiGates in the virtual cluster. If one of the cluster units fails, the remaining unit will process traffic for both VDOMs.

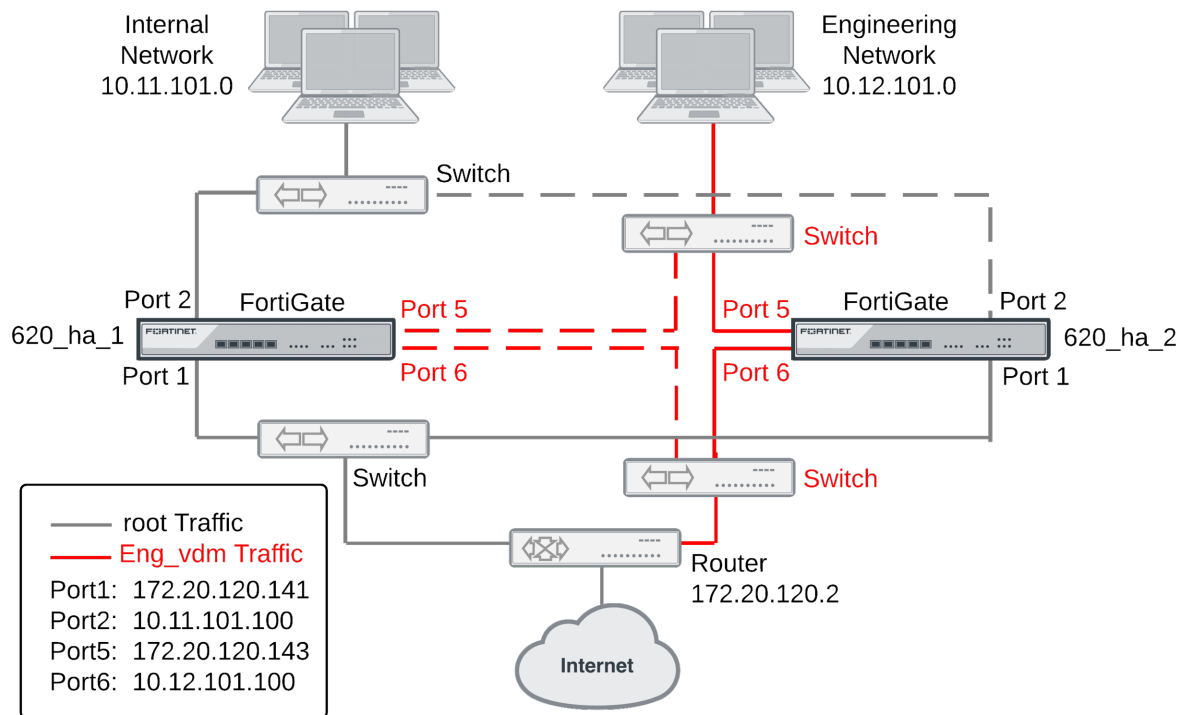
The procedures in this example describe some of many possible sequences of steps for configuring virtual clustering. For simplicity many of these procedures assume that you are starting with new FortiGates set to the factory default configuration. However, this is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

Example virtual clustering network topology

The following figure shows a typical FortiGate HA virtual cluster consisting of two FortiGates (FGT_ha_1 and FGT_ha_2) connected to an internal network, an engineering network and the Internet. To simplify the diagram the heartbeat connections are not shown.

The traffic from the internal network is processed by the root VDOM, which includes the port1 and port2 interfaces. The traffic from the engineering network is processed by the Eng_vdm VDOM, which includes the port5 and port6 interfaces. VDOM partitioning is configured so that all traffic from the internal network is processed by FGT_ha_1 and all traffic from the engineering network is processed by FGT_ha_2.

This virtual cluster uses the default FortiGate heartbeat interfaces (port3 and port4).

Example virtual cluster showing VDOM partitioning**General configuration steps**

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. Configure VDOM settings for the cluster:
 - Enable multiple VDOMs.
 - Add the Eng_vdm VDOM.
 - Add port5 and port6 to the Eng_vdm.
4. Configure VDOM partitioning.
5. Confirm that the cluster units are operating as a virtual cluster and add basic configuration settings to the cluster.
 - View cluster status from the GUI or CLI.
 - Add a password for the admin administrative account.

- Change the IP addresses and netmasks of the port1, port2, port5, and port6 interfaces.
- Add a default routes to each VDOM.

Configuring virtual clustering with two VDOMs and VDOM partitioning - GUI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	<div> ✓ Registered (bdickie@fortinet.com) Launch Portal </div>
	FortiGuard	IPS & Application Control	<div> ✓ Licensed (Expires 2016-08-22) </div>
		AntiVirus	<div> ✓ Licensed (Expires 2016-08-22) </div>
		Web Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
		Anti-Spam Filtering	<div> ✓ Licensed (Expires 2016-08-21) </div>
	FortiCloud	Account	<div> Activate </div>
	FortiSandbox	FortiSandbox Appliance	<div> ✗ Not Configured Configure </div>
	FortiClient	Status	<div> ✓ Free License <div> How to Purchase </div> </div>
		Clients Registered	<div> <input type="text"/> 0 of 10 Enter License </div>
		FortiClient Installers	<div> Details </div>
	FortiToken Mobile	Tokens Assigned	<div> <input type="text"/> 0 of 2 </div>

2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
4. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
----------	----------

5. Select **OK**.
6. Go to **System > HA** and change the following settings.

Mode	Active-Passive
------	----------------

Group Name	vexample2.com
Password	vHA_pass_2

7. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces (see [Cluster virtual MAC addresses on page 224](#)). The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate.
8. Repeat these steps for the second FortiGate.
Set the second FortiGate host name to:

New Name	FGT_ha_2
-----------------	----------

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 to switch connected to the Internet.
You could use the same switch for the port1 and port5 interfaces.
3. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
4. Connect the port6 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the engineering network.
5. Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
6. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
7. Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete you can continue.

To configure VDOM settings for the cluster

1. Log into the GUI.
2. On the **System Information** dashboard widget, beside **Virtual Domain** select **Enable**.
3. Select **OK** and then log back into the GUI.
4. Go to **System > VDOM** and select **Create New** to add a new VDOM.

Virtual Domain	Eng_vdm
-----------------------	---------

5. Go to **Network > Interfaces**.
6. Edit the **port5** interface, add it to the Eng_vdm VDOM and configure other interface settings:

Alias	Engineering_external
Virtual Domain	Eng_vdm
IP/Netmask	172.20.120.143/24

7. Select **OK**.
8. Edit the **port6** interface, add it to the Eng_vdm VDOM and configure other interface settings:

Alias	Engineering_internal
--------------	----------------------

Virtual Domain	Eng_vdm
IP/Netmask	10.120.101.100/24
Administrative Access	HTTPS, PING, SSH

9. Select **OK**.

To add a default route to each VDOM

1. Use the VDOM meny to enter the root VDOM.
2. Go to **Network > Static Routes**.
3. Change the default route.

Destination	0.0.0.0/0.0.0.0
Device	port1
Gateway	172.20.120.2
Administrative Distance	10

4. Select **Global**.
5. Enter the Eng_vdm VDOM.
6. Go to **Network > Static Routes**.
7. Change the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port5
Gateway	172.20.120.2
Distance	10

To configure VDOM partitioning

1. Go to **System > HA**.
The cluster members shows two cluster units in Virtual Cluster 1.
2. Edit the cluster unit with the **Role** of **MASTER**.
3. Change **VDOM partitioning** to move the **Eng_vdm** to the **Virtual Cluster 2** list.
4. Select **OK**.
5. Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit to the following:

Device Priority		
Host Name	Virtual Cluster 1	Virtual Cluster 2

FGT_ha_1	200	100
FGT_ha_2	100	200

You can do this by editing the HA configurations of each cluster unit in the cluster members list and changing device priorities.

Since the device priority of Virtual Cluster 1 is highest for FGT_ha_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by FGT_ha_1.

Since the device priority of Virtual Cluster 2 is highest for FGT_ha_2 and since the Eng_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng_vdm VDOM is processed by FGT_ha_2.

To view cluster status and verify the VDOM partitioning configuration

1. Log into the GUI.
2. Go to **System > HA**.

The cluster members list should show the following:

- Virtual Cluster 1 contains the root VDOM.
- FGT_ha_1 is the primary unit (master) for Virtual Cluster 1.
- Virtual Cluster 2 contains the Eng_vdm VDOM.
- FGT_ha_2 is the primary unit (master) for Virtual Cluster 2.

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for the root VDOM is processed by FGT_ha_1 and traffic for the Eng_vdm is processed by FGT_ha_2.

1. Log into the GUI by connecting to port2 using IP address 10.11.101.100.
You will log into FGT_ha_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by FGT_ha_1. You can confirm that you have logged into FGT_ha_1 by checking the host name on the System Information dashboard widget.
2. Log into the GUI by connecting to port6 using IP address 10.12.101.100.
You will log into FGT_ha_2 because port6 is in the Eng_vdm VDOM and all traffic for this VDOM is processed by FGT_ha_2.
3. Add security policies to the root virtual domain that allows communication from the internal network to the Internet and connect to the Internet from the internal network.
4. Log into the GUI and go to **System > HA** and select **View HA Statistics**.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the FGT_ha_1 unit.
5. Add security policies to the Eng_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.
6. Log into the GUI and go to **System > HA** and select **View HA Statistics**.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the FGT_ha_2 unit.

Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. Change the host name for this FortiGate:

```
config system global
    set hostname FGT_ha_1
end
```

4. Configure HA settings.

```
config system ha
    set mode a-p
    set group-name vexample2.com
    set password vHA_pass_2
end
```

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces (see [Cluster virtual MAC addresses on page 224](#)). The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11

- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Power off the first FortiGate.
5. Repeat these steps for the second FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the Internet.
2. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 to switch connected to the Internet.
You could use the same switch for port1 and port5.
3. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
4. Connect the port6 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the engineering network.
5. Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
6. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
7. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete you can continue.

To configure VDOM settings for the cluster

1. Log into the CLI.
2. Enter the following command to enable multiple VDOMs for the cluster.

```
config system global
    set vdom-admin enable
```

```
end
```

3. Log back into the CLI.

4. Enter the following command to add the Eng_vdm VDOM:

```
config vdom
  edit Eng_vdm
end
```

5. Edit the port5 interface, add it to the Eng_vdm VDOM and configure other interface settings:

```
config global
  config system interface
    edit port5
      set vdom Eng_vdm
      set alias Engineering_external
      set ip 172.20.12.143/24
    next
    edit port6
      set vdom Eng_vdm
      set alias Engineering_internal
      set ip 10.120.101.100/24
    end
  end
end
```

To add a default route to each VDOM

1. Enter the following command to add default routes to the root and Eng_vdm VDOMs.

```
config vdom
  edit root
    config router static
      edit 1
        set dst 0.0.0.0/0.0.0.0
        set gateway 172.20.120.2
        set device port1
      end
    next
    edit Eng_vdm
      config router static
        edit 1
          set dst 0.0.0.0/0.0.0.0
          set gateway 172.20.120.2
          set device port5
        end
      end
    end
end
```

To configure VDOM partitioning

1. Enter the `get system ha status` command to view cluster unit status:

For example, from the FGT_ha_2 cluster unit CLI:

```
config global
  get system ha status
  .
  .
  .
  number of vcluster: 1
  vcluster 1: work 169.254.0.1
```

```
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

This command output shows that VDOM partitioning has not been configured because only virtual cluster 1 is shown. The command output also shows that the FGT_ha_2 is the primary unit for the cluster and for virtual cluster 1 because this cluster unit has the highest serial number

2. Enter the following commands to configure VDOM partitioning:

```
config global
  config system ha
    set vcluster2 enable
    config secondary-vcluster
      set vdom Eng_vdm
    end
  end
end
```

3. Enter the `get system ha status` command to view cluster unit status:

For example, from the FGT_ha_2 cluster unit CLI:

```
config global
  get system ha status
  .
  .
  .
  number of vcluster: 2
  vcluster 1: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705
  vcluster 2: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705
```

This command output shows VDOM partitioning has been configured because both virtual cluster 1 and virtual cluster 2 are visible. However the configuration is not complete because FGT_ha_2 is the primary unit for both virtual clusters. The command output shows this because under both vcluster entries the Master entry shows FG600B3908600825, which is the serial number of FGT_ha_2. As a result of this configuration, FGT_ha_2 processes traffic for both VDOMs and FGT_ha_1 does not process any traffic.

4. Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit so that FGT_ha_1 processes virtual cluster 1 traffic and FGT_ha_2 processes virtual cluster 2 traffic.

Since the root VDOM is in virtual cluster 1 and the Eng_vdm VDOM is in virtual cluster 2 the result of this configuration will be that FGT_ha_1 will process all root VDOM traffic and FGT_ha_2 will process all Eng_vdm traffic. You make this happen by changing the cluster unit device priorities for each virtual cluster. You could use the following settings:

Host Name	Device Priority	
	Virtual Cluster 1	Virtual Cluster 2
FGT_ha_1	200	100
FGT_ha_2	100	200

Since the device priority is not synchronized you can edit the device priorities of each virtual cluster on

each FortiGate separately. To do this:

- Log into the CLI and note the FortiGate you have actually logged into (for example, by checking the host name displayed in the CLI prompt).
- Change the virtual cluster 1 and 2 device priorities for this cluster unit.
- Then use the `execute ha manage` command to log into the other cluster unit CLI and set its virtual cluster 1 and 2 device priorities.

Enter the following commands from the FGT_ha_1 cluster unit CLI:

```
config global
  config system ha
    set priority 200
    config secondary-vcluster
      set priority 100
    end
  end
end
```

Enter the following commands from the FGT_ha_2 cluster unit CLI:

```
config global
  config system ha
    set priority 100
    config secondary-vcluster
      set priority 200
    end
  end
end
```



The cluster may renegotiate during this step resulting in a temporary loss of connection to the CLI and a temporary service interruption.

Since the device priority of Virtual Cluster 1 is highest for FGT_ha_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by FGT_ha_1.

Since the device priority of Virtual Cluster 2 is highest for FGT_ha_2 and since the Eng_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng_vdm VDOM is processed by FGT_ha_2.

To verify the VDOM partitioning configuration

1. Log into the FGT_ha_2 cluster unit CLI and enter the following command:

```
config global
  get system ha status
  .
  .
  .
  number of vcluster: 2
  vcluster 1: standby 169.254.0.2
  Slave :1 FG600B3908600825
  Master:0 FG600B3908600705
  vcluster 2: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705
```

The command output shows that FGT_ha_1 is the primary unit for virtual cluster 1 (because the

command output shows the `Master` of virtual cluster 1 is the serial number of FGT_ha_1) and that FGT_ha_2 is the primary unit for virtual cluster 2.

If you enter the same command from the FGT_ha_1 CLI the same information is displayed but in a different order. The command always displays the status of the cluster unit that you are logged into first.

```
config global
get system ha status
.
.
.
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:0 FG600B3908600705
Slave :1 FG600B3908600825
vcluster 2: standby 169.254.0.1
Slave :1 FG600B3908600705
Master:0 FG600B3908600825
```

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for the root VDOM is processed by FGT_ha_1 and traffic for the Eng_vdm is processed by FGT_ha_2. These steps assume the cluster is operating correctly.

1. Log into the CLI by connecting to port2 using IP address 10.11.101.100.
You will log into FGT_ha_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by FGT_ha_1. You can confirm that you have logged into FGT_ha_1 by checking the host name in the CLI prompt. Also the `get system status` command displays the status of the FGT_ha_1 cluster unit.
2. Log into the GUI or CLI by connecting to port6 using IP address 10.12.101.100.
You will log into FGT_ha_2 because port6 is in the Eng_vdm VDOM and all traffic for this VDOM is processed by FGT_ha_2.
3. Add security policies to the root virtual domain that allow communication from the internal network to the Internet and connect to the Internet from the internal network.
4. Log into the GUI and go to **System > HA > View HA Statistics**.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the FGT_ha_1 unit.
5. Add security policies to the Eng_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.
6. Log into the GUI and go to **System > HA > View HA Statistics**.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the FGT_ha_2 unit.

Example inter-VDOM links in a virtual clustering configuration

In a virtual domain configuration you can use inter-VDOM links to route traffic between two virtual domains operating in a single FortiGate without using physical interfaces. Adding an inter-VDOM link has the affect of

adding two interfaces to the FortiGate and routing traffic between the virtual domains using the inter-VDOM link interfaces.

In a virtual clustering configuration inter-VDOM links can only be made between virtual domains that are in the same virtual cluster. So, if you are planning on configuring inter-VDOM links in a virtual clustering configuration, you should make sure the virtual domains that you want to link are in the same virtual cluster.

For example, the following tables show an example virtual clustering configuration where each virtual cluster contains four virtual domains. In this configuration you can configure inter-VDOM links between root and vdom_1 and between vdom_2 and vdom_3. But, you cannot configure inter-VDOM links between root and vdom_2 or between vdom_1 and vdom_3 (and so on).

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
root	Priority	Priority
	200	100
vdom_1	Role	Role
	Primary	Subordinate

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
vdom_2	Priority	Priority
	100	200
vdom_3	Role	Role
	Subordinate	Primary

Configuring inter-VDOM links in a virtual clustering configuration

Configuring inter-VDOM links in a virtual clustering configuration is very similar to configuring inter-VDOM links for a standalone FortiGate. The main difference the `config system vdom-link` command includes the `vcluster` keyword. The default setting for `vcluster` is `vcluster1`. So you only have to use the `vcluster` keyword if you are adding an inter-VDOM link to virtual cluster 2.

To add an inter-VDOM link to virtual cluster 1

This procedure describes how to create an inter-VDOM link to virtual cluster 1 that results in a link between the root and vdom_1 virtual domains.



Inter-VDOM links are also called internal point-to-point interfaces.

1. Add an inter-VDOM link called `vc1link`.

```
config global
  config system vdom-link
    edit vc1link
  end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc1link0` and `vc1link1`. These interfaces appear in all CLI and GUI interface lists. These interfaces can only be added to virtual domains in virtual cluster 1.

2. Bind the `vc1link0` interface to the root virtual domain and bind the `vc1link1` interface to the `vdom_1` virtual domain.

```
config system interface
  edit vc1link0
    set vdom root
  next
  edit vc1link1
    set vdom vdom_1
  end
```

To add an inter-VDOM link to virtual cluster 2

This procedure describes how to create an inter-VDOM link to virtual cluster 2 that results in a link between the `vdom_2` and `vdom_3` virtual domains.

1. Add an inter-VDOM link called `vc2link`.

```
config global
  config system vdom-link
    edit vc2link
      set vcluster vcluster2
    end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc2link0` and `vc2link1`. These interfaces appear in all CLI and GUI interface lists. These interfaces can only be added to virtual domains in virtual cluster 2.

2. Bind the `vc2link0` interface to the `vdom_2` virtual domain and bind the `vc2link1` interface to the `vdom_3` virtual domain.

```
config system interface
  edit vc2link0
    set vdom vdom_2
  next
  edit vc2link1
    set vdom vdom_3
  end
```

Troubleshooting virtual clustering

Troubleshooting virtual clusters is similar to troubleshooting any cluster (see [FGCP configuration examples and troubleshooting on page 66](#)). This section describes a few testing and troubleshooting techniques for virtual clustering.

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for different VDOMs will be distributed among both FortiGates in the virtual cluster. These steps assume the cluster is otherwise operating correctly.

1. Log into the GUI or CLI using the IP addresses of interfaces in each VDOM.
Confirm that you have logged into the FortiGate that should be processing traffic for that VDOM by checking the HTML title displayed by your web browser or the CLI prompt. Both of these should include the host name of the cluster unit that you have logged into. Also on the system Dashboard, the System Information widget displays the serial number of the FortiGate that you logged into. From the CLI the `get system status` command displays the status of the cluster unit that you logged into.
2. To verify that the correct cluster unit is processing traffic for a VDOM:
 - Add security policies to the VDOM that allow communication between the interfaces in the VDOM.
 - Optionally enable traffic logging and other monitoring for that VDOM and these security policies.
 - Start communication sessions that pass traffic through the VDOM.
 - Log into the GUI and go to **System > HA > View HA Statistics**. Verify that the statistics display shows more active sessions, total packets, network utilization, and total bytes for the unit that should be processing all traffic for the VDOM.
 - Optionally check traffic logging and the Top Sessions Widget for the FortiGate that should be processing traffic for that VDOM to verify that the traffic is being processed by this FortiGate.

Full mesh HA

This chapter provides an introduction to full mesh HA and also contains general procedures and configuration examples that describe how to configure FortiGate full mesh HA.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

Full mesh HA overview

When two or more FortiGates are connected to a network in an HA cluster the reliability of the network is improved because the HA cluster replaces a single FortiGate as a single point of failure. With a cluster, a single FortiGate is replaced by a cluster of two or more FortiGates.

However, even with a cluster, potential single points of failure remain. The interfaces of each cluster unit connect to a single switch and that switch provides a single connection to the network. If the switch fails or if the connection between the switch and the network fails service is interrupted to that network.

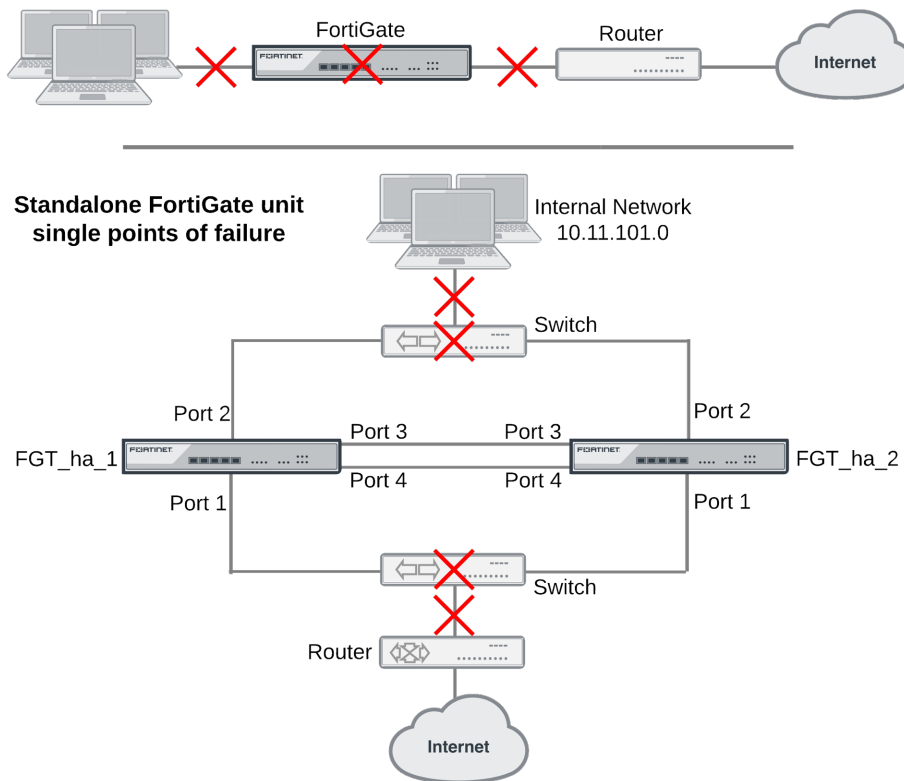
The HA cluster does improve the reliability of the network because switches are not as complex components as FortiGates, so are less likely to fail. However, for even greater reliability, a configuration is required that includes redundant connections between the cluster the networks that it is connected to.

FortiGate models that support 802.3ad Aggregate or Redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster.

This redundant configuration can be achieved using FortiGate 802.3ad Aggregate or Redundant interfaces and a full mesh HA configuration. In a full mesh HA configuration, you connect an HA cluster consisting of two or more FortiGates to the network using 802.3ad Aggregate or Redundant interfaces and redundant switches. Each 802.3ad Aggregate or Redundant interface is connected to two switches and both of these switches are connected to the network. In addition you must set up an IEEE 802.1Q (also called Dot1Q) or ISL link between the redundant switches connected to the Aggregate or Redundant interfaces.

The resulting full mesh configuration, an example is shown below, includes redundant connections between all network components. If any single component or any single connection fails, traffic automatically switches to the redundant component and connection and traffic flow resumes.

Single points of failure in a standalone and HA network configuration



Full mesh HA and redundant heartbeat interfaces

A full mesh HA configuration also includes redundant HA heartbeat interfaces. At least two heartbeat interfaces should be selected in the HA configuration and both sets of HA heartbeat interfaces should be connected. The HA heartbeat interfaces do not have to be configured as redundant interfaces because the FGCP handles failover between heartbeat interfaces.

Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces

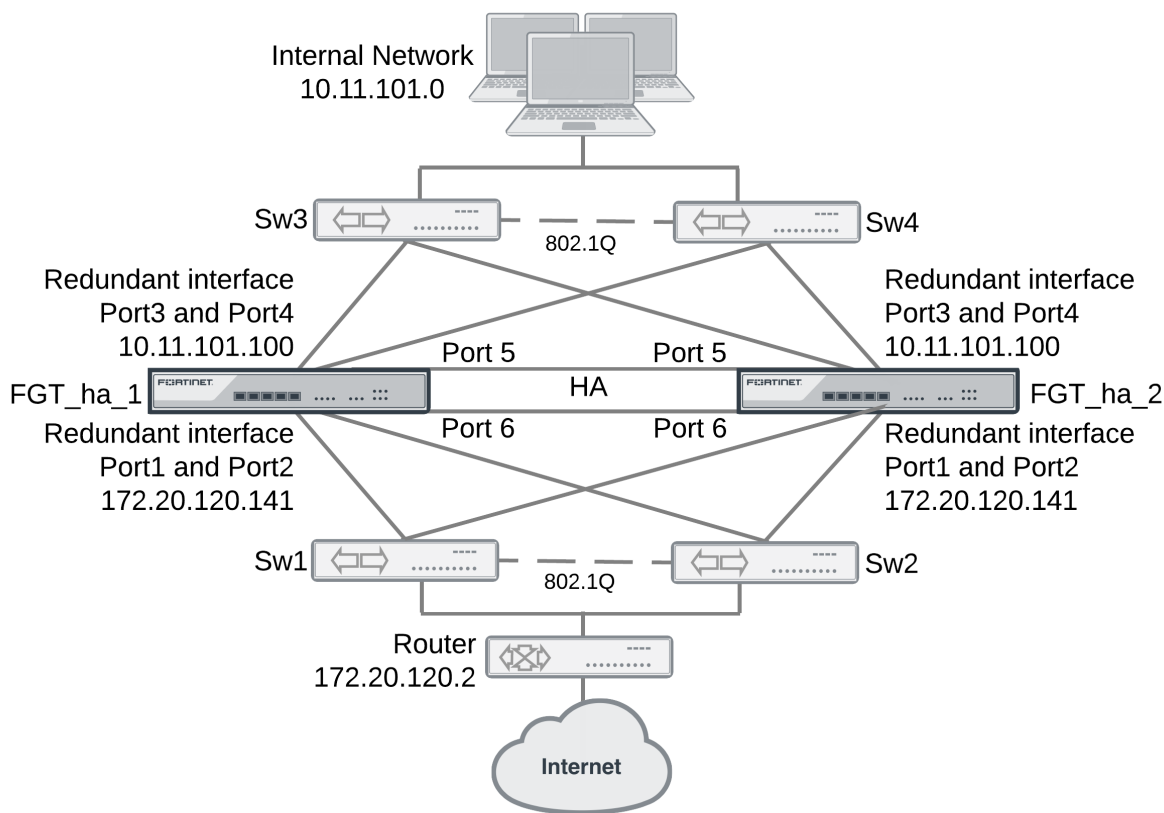
Full mesh HA is supported for both redundant interfaces and 802.3ad aggregate interfaces. In most cases you would simply use redundant interfaces. However, if your switches support 802.3ad aggregate interfaces and split multi-trunking you can use aggregate interfaces in place of redundant interfaces for full mesh HA. One advantage of using aggregate interfaces is that all of the physical interfaces in the aggregate interface can send and receive packets. As a result, using aggregate interfaces may increase the bandwidth capacity of the cluster.

Usually redundant and aggregate interfaces consist of two physical interfaces. However, you can add more than two physical interfaces to a redundant or aggregate interface. Adding more interfaces can increase redundancy protection. Adding more interfaces can also increase bandwidth capacity if you are using 802.3ad aggregate interfaces.

Example full mesh HA configuration

The following figure shows a full mesh HA configuration with a cluster of two FortiGate. This section describes the FortiGate configuration settings and network components required for a full mesh HA configuration. This section also contains example steps for setting up this full mesh HA configuration. The procedures in this section describe one of many possible sequences of steps for configuring full mesh HA. As you become more experienced with FortiOS, HA, and full mesh HA you may choose to use a different sequence of configuration steps.

Full Mesh HA configuration



For simplicity these procedures assume that you are starting with two new FortiGates set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

These procedures describe how to configure a cluster operating in NAT/Route mode because NAT/Route is the default FortiGate operating mode. However, the steps are the same if the cluster operates in Transparent mode. You can either switch the cluster units to operate in Transparent mode before beginning these procedures, or you can switch the cluster to operate in Transparent mode after HA is configured and the cluster is connected and operating.

Full mesh HA configuration

The two FortiGates (FGT_ha_1 and FGT_ha_2) can be operating in NAT/Route or Transparent mode. Aside from the standard HA settings, the FortiGate configuration includes the following:

- The port5 and port6 interfaces configured as heartbeat interfaces. A full mesh HA configuration also includes redundant HA heartbeat interfaces.
- The port1 and port2 interfaces added to a redundant interface. Port1 is the active physical interface in this redundant interface. To make the port1 interface the active physical interface it should appear above the port2 interface in the redundant interface configuration.
- The port3 and port4 interfaces added to a redundant interface. Port3 is the active physical interface in this redundant interface. To make the port3 interface the active physical interface it should appear above the port4 interface in the redundant interface configuration.

Full mesh switch configuration

The following redundant switch configuration is required:

- Two redundant switches (Sw3 and Sw4) connected to the internal network. Establish an 802.1Q (Dot1Q) or interswitch-link (ISL) connection between them.
- Two redundant switches (Sw1 and Sw2) connected to the Internet. Establish an 802.1Q (Dot1Q) or interswitch-link (ISL) connection between them.

Full mesh network connections

Make the following physical network connections for FGT_ha_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

Make the following physical network connections for FGT_ha_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

How packets travel from the internal network through the full mesh cluster and to the Internet

If the cluster is operating in active-passive mode and FGT_ha_2 is the primary unit, all packets take the following path from the internal network to the internet:

1. From the internal network to Sw4. Sw4 is the active connection to FGT_ha_2; which is the primary unit. The primary unit receives all packets.
2. From Sw4 to the FGT_ha_2 port3 interface. Active connection between Sw4 and FGT_ha_2. Port3 is the active member of the redundant interface.

3. From FGT_ha_2 port3 to FGT_ha_2 port1. Active connection between FGT_ha_2 and Sw2. Port1 is the active member of the redundant interface.
4. From Sw2 to the external router and the Internet.

Configuring full-mesh HA - GUI

Each cluster unit must have the same HA configuration.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing and entering a license key if you purchased more than 10 **Virtual Domains (VDOMS)**. All of the FortiGates in a cluster must have the same level of licensing.

License Information			
	Support Contract	Registration	Registered (bdickie@fortinet.com) Launch Portal
	FortiGuard	IPS & Application Control	Licensed (Expires 2016-08-22)
		AntiVirus	Licensed (Expires 2016-08-22)
		Web Filtering	Licensed (Expires 2016-08-21)
		Anti-Spam Filtering	Licensed (Expires 2016-08-21)
	FortiCloud	Account	Activate
	FortiSandbox	FortiSandbox Appliance	Not Configured Configure
	FortiClient	Status	Free License How to Purchase Enter License
		Clients Registered	0 of 10 Details
		FortiClient Installers	
	FortiToken Mobile	Tokens Assigned	0 of 2

2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate. We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
4. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

5. Go to **System > HA** and change the following settings.

Mode	Active-Active
-------------	---------------

Group Name	Rexample1.com	
Password	RHA_pass_1	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

6. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
```



```

.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.

```

6. Power off the first FortiGate.
7. Repeat these steps for the second FortiGate.
Set the second FortiGate host name to:

New Name	FGT_ha_2
-----------------	----------

To connect the cluster to your network

1. Make the following physical network connections for FGT_ha_1:
 - Port1 to Sw1 (active)
 - Port2 to Sw2 (inactive)
 - Port3 to Sw3 (active)
 - Port4 to Sw4 (inactive)
2. Make the following physical network connections for FGT_ha_2:
 - Port1 to Sw2 (active)
 - Port2 to Sw1 (inactive)
 - Port3 to Sw4 (active)
 - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable 802.1Q (Dot1Q) or ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.
The System Information dashboard widget shows the **Cluster Name** (Rexample1.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
2. Go to **System > HA** to view the cluster members list.
The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGates are not functioning as a cluster. See [Example full mesh HA configuration on page 165](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Admin > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Router > Static > Static Routes** and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to **System > Network > Interfaces** and select **Create New** and configure the redundant interface to connect to the Internet.

Name	Port1_Port2
Type	Redundant
Physical Interface Members	
Selected Interfaces	port1, port2
IP/Netmask	172.20.120.141/24

8. Select **OK**.
9. Select **Create New** and configure the redundant interface to connect to the internal network.

Name	Port3_Port4
Type	Redundant
Physical Interface Members	
Selected Interfaces	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

10. Select **OK**.
The virtual MAC addresses of the FortiGate interfaces change to the following. Notice that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

11. Go to **Router > Static > Static Routes**.

12. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

13. Select **OK**.

To configure HA port monitoring for the redundant interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the redundant interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4. Select **OK**.

Configuring Full Mesh HA - CLI

Each cluster must have the same HA configuration. Use the following procedure to configure the FortiGates for HA operation.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
3. Enter a new Host Name for this FortiGate.

```
config system global
    set hostname FGT_ha_1
end
```

4. Configure HA settings.

```
config system ha
    set mode a-a
    set group-name Rexample1.com
    set password RHA_pass_1
    set hbdev port5 50 port6 50
end
```

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11

- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Power off the first FortiGate.
5. Repeat these steps for the second FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to your network

1. Make the following physical network connections for FGT_ha_1:
 - Port1 to Sw1 (active)
 - Port2 to Sw2 (inactive)
 - Port3 to Sw3 (active)
 - Port4 to Sw4 (inactive)
2. Make the following physical network connections for FGT_ha_2:
 - Port1 to Sw2 (active)
 - Port2 to Sw1 (inactive)
 - Port3 to Sw4 (active)
 - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable 802.1Q (Dot1Q) or ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into.
 If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit.
 If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit.
 If the command output includes `Current HA mode: standalone` the cluster unit is not operating in HA mode.
3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

4. Use the `execute ha manage` command to connect to the other cluster unit's CLI and use these commands to verify cluster status.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGates are not functioning as a cluster. See [Example full mesh HA configuration on page 165](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings. Some steps use the CLI and some the GUI.

1. Log into the cluster CLI.
2. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <password_str>
  end
```

3. Temporarily delete the default route.
 You cannot add an interface to a redundant interface if any settings (such as the default route) are

configured for it.

```
config router static
  delete 1
end
```

4. Go to **System > Network > Interface** and select **Create New** to add the redundant interface to connect to the Internet.
5. Add the redundant interface to connect to the Internet.

```
config system interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
  end
```

6. Add the redundant interface to connect to the internal network.

```
config system interface
  edit Port3_Port4
    set type redundant
    set member port3 port4
  end
```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

7. Go to **Router > Static > Static Routes**.
8. Add the default route.

```
config router static
  edit 1
```

```
set dst 0.0.0.0 0.0.0.0
set gateway 172.20.120.2
set device Port1_Port2
end
```

To configure HA port monitoring for the redundant interfaces

1. Enter the following command to configure port monitoring for the redundant interfaces:

```
config system ha
set monitor Port1_Port2 Port3_Port4
end
```

Troubleshooting full mesh HA

Troubleshooting full mesh HA clusters is similar to troubleshooting any cluster (see [FGCP configuration examples and troubleshooting on page 66](#) or [Virtual clusters on page 143](#)). The configuration and operation of a full mesh HA cluster is very similar to the configuration and operation of a standard cluster. The only differences relate to the configuration, connection, and operation of the redundant interfaces and redundant switches.

- Make sure the redundant interfaces and switches are connected correctly. With so many connections it is possible to make mistakes or for cables to become disconnected.
- Confirm that the configuration of the cluster unit 802.3ad Aggregate or Redundant interfaces is correct according to the configuration procedures in this chapter.
- In some configurations with some switch hardware, MAC-learning delays on the inter-switch links on the surrounding topologies may occur. The delays occur if the gratuitous ARP packets sent by the cluster after a failover are delayed by the switches before being sent across the inter-switch link. If this happens the surrounding topologies may be delayed in recognizing the failover and will keep sending packets to the MAC address of the failed primary unit resulting in lost traffic. Resolving this problem may require changing the configuration of the switch or replacing them with switch hardware that does not delay the gratuitous ARP packets.

Operating clusters and virtual clusters

With some exceptions, you can operate a cluster in much the same way as you operate a standalone FortiGate. This chapter describes those exceptions and also the similarities involved in operating a cluster instead of a standalone FortiGate.

Operating a cluster

The configurations of all of the FortiGates in a cluster are synchronized so that the cluster units can simulate a single FortiGate. Because of this synchronization, you manage the HA cluster instead of managing the individual cluster units. You manage the cluster by connecting to the GUI using any cluster interface configured for HTTPS or HTTP administrative access. You can also manage the cluster by connecting to the CLI using any cluster interface configured for SSH or telnet administrative access.

The cluster GUI dashboard displays the cluster name, the host name and serial number of each cluster member, and also shows the role of each unit in the cluster. The roles can be master (primary unit) and slave (subordinate units). The dashboard also displays a cluster unit front panel illustration.

You can also go to **System > HA** to view the cluster members list. This includes status information for each cluster unit. You can also use the cluster members list for a number of cluster management functions including changing the HA configuration of an operating cluster, changing the host name and device priority of a subordinate unit, and disconnecting a cluster unit from a cluster. See [Cluster members list on page 194](#).

You can use log messages to view information about the status of the cluster. See [Clusters and logging on page 185](#).

You can use SNMP to manage the cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration information and receive traps. See [Clusters and SNMP on page 188](#).

You can configure a reserved management interface to manage individual cluster units. You can use this interface to access the GUI or CLI and to configure SNMP management for individual cluster units. See [Managing individual cluster units using a reserved management interface on page 178](#).

You can manage individual cluster units by using SSH, telnet, or the CLI console on the GUI dashboard to connect to the CLI of the cluster. From the CLI you can use the `execute ha manage` command to connect to the CLI of any unit in the cluster.

You can also manage individual cluster units by using a null-modem cable to connect to any cluster unit CLI. From there you can use the `execute ha manage` command to connect to the CLI of each unit in the cluster.

Operating a virtual cluster

Managing a virtual cluster is very similar to managing a cluster that does not contain multiple virtual domains. Most of the information in this chapter applies to managing both kinds of clusters. This section describes what is different when managing a virtual cluster.

If virtual domains are enabled, the cluster GUI dashboard displays the cluster name and the role of each cluster unit in virtual cluster 1 and virtual cluster 2.

The configuration and maintenance options that you have when you connect to a virtual cluster GUI or CLI depend on the virtual domain that you connect to and the administrator account that you use to connect.

If you connect to a cluster as the administrator of a virtual domain, you connect directly to the virtual domain. Since HA virtual clustering is a global configuration, virtual domain administrators cannot see HA configuration options. However, virtual domain administrators see the host name of the cluster unit that they are connecting to on the web browser title bar or CLI prompt. This host name is the host name of the primary unit for the virtual domain. Also, when viewing log messages the virtual domain administrator can select to view log messages for either of the cluster units.

If you connect to a virtual cluster as the admin administrator you connect to the global GUI or CLI. Even so, you are connecting to an interface and to the virtual domain that the interface has been added to. The virtual domain that you connect to does not make a difference for most configuration and maintenance operations. However, there are a few exceptions. You connect to the FortiGate that functions as the primary unit for the virtual domain. So the host name displayed on the web browser title bar and on the CLI is the host name of this primary unit.

Managing individual cluster units using a reserved management interface

You can provide direct management access to all cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. Configuration changes to the reserved management interface are not synchronized to other cluster units.

The reserved management interface provides direct management access to each cluster unit and gives each cluster unit a different identity on your network. This simplifies using external services, such as SNMP, to separately monitor and manage each cluster unit.



The reserved management interface is not assigned an HA virtual MAC address like other cluster interfaces. Instead the reserved management interface retains the permanent hardware address of the physical interface unless you change it using the `config system interface` command.

The reserved management interface and IP address should not be used for managing a cluster using FortiManager. To correctly manage a FortiGate HA cluster with FortiManager use the IP address of one of the cluster unit interfaces.

If you enable SNMP administrative access for the reserved management interface you can use SNMP to monitor each cluster unit using the reserved management interface IP address. To monitor each cluster unit using SNMP, just add the IP address of each cluster unit's reserved management interface to the SNMP server configuration. You must also enable direct management of cluster members in the cluster SNMP configuration.

If you enable HTTPS or HTTP administrative access for the reserved management interfaces you can connect to the GUI of each cluster unit. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. From the subordinate units the GUI has the same features as the primary unit except that unit-specific information is displayed for the subordinate unit, for example:

- The **Dashboard System Information** widget displays the subordinate unit serial number but also displays the same information about the cluster as the primary unit
- On the Cluster members list (go to **System > HA**) you can change the HA configuration of the subordinate unit that you are logged into. For the primary unit and other subordinate units you can change only the host name and device priority.
- Log Access displays the logs of the subordinate that you are logged into first, You use the HA Cluster list to view the log messages of other cluster units including the primary unit.

If you enable SSH or TELNET administrative access for the reserved management interfaces you can connect to the CLI of each cluster unit. The CLI prompt contains the host name of the cluster unit that you have connected to. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. You can also use the `execute ha manage` command to connect to other cluster unit CLIs.

The reserved management interface is available in NAT/Route and in Transparent mode. It is also available if the cluster is operating with multiple VDOMs. In Transparent mode you cannot normally add an IP address to an interface. However, you can add an IP address to the reserved management interface.

Using the HA reserved management interface for FortiSandbox, SNMP and other management services

By default, management services such as SNMP, remote logging, remote authentication and communication with FortiSandbox and so on use a cluster interface. As a result communication from each cluster unit comes from a cluster interface instead of from the interface of an individual cluster unit and not from the HA reserved management interface.

If you want to use the HA reserved management interface for these features you must enter the following command:

```
config system ha
    set ha-direct enable
end
```

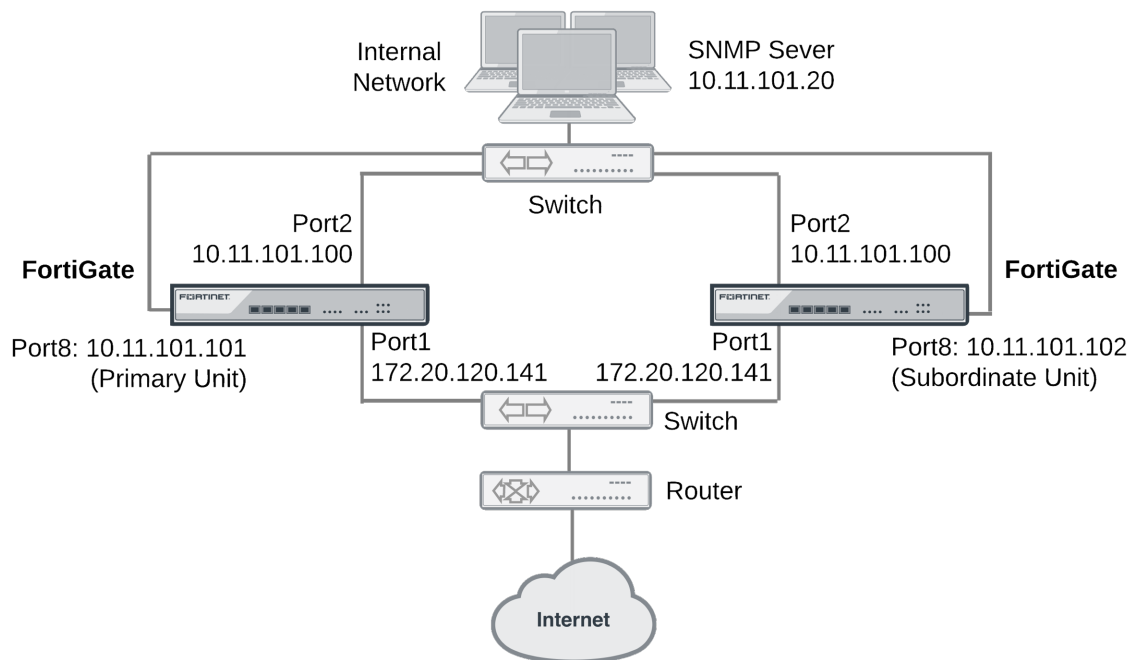
The result is that all management services use the HA reserved management interface. This means that individual cluster units send log messages and communicate with FortiSandbox and so on using the HA reserved management interface instead of one of the cluster interfaces. This allows you to manage each cluster unit separately and to separate the management traffic from each cluster unit. This can also be useful if each cluster unit is in a different location.

If you just want to use the HA reserved management interface for SNMP remote management you can enable `ha-direct` in the SNMP configuration as shown in the following example.

Configuring the reserved management interface and SNMP remote management of individual cluster units

This example describes how to configure SNMP remote management of individual cluster units using the HA reserved management interface. The configuration consists of two FortiGate-620B units already operating as a cluster. In the example, the port8 interface of each cluster unit is connected to the internal network using the switch and configured as the reserved management interface.

SNMP remote management of individual cluster units



To configure the reserved management interface - GUI

1. Go to **System > HA**.
2. Edit the primary unit.
3. Select **Reserve Management Port for Cluster Member** and select port8.
4. Select **OK**.

To configure the reserved management interface - CLI

From the CLI you can also configure IPv4 and IPv6 default routes that are only used by the reserved management interface.

1. Log into the CLI of any cluster unit.
2. Enter the following command to enable the reserved management interface, set port8 as the reserved interface, and add an IPv4 default route of 10.11.101.2 and an IPv6 default route of 2001:db8:0:2::20 for the reserved management interface.

```
config system ha
    set ha-mgmt-status enable
    set ha-mgmt-interface port8
    set ha-mgmt-interface-gateway 10.11.101.2
    set ha-mgmt-interface-gateway6 2001:db8:0:2::20
end
```

The reserved management interface default route is not synchronized to other cluster units.

To change the primary unit reserved management interface configuration - GUI

You can change the IP address of the primary unit reserved management interface from the primary unit GUI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. From a PC on the internal network, browse to `http://10.11.101.100` and log into the cluster GUI.

This logs you into the primary unit GUI.

You can identify the primary unit from its serial number or host name that appears on the System Information dashboard widget.

2. Go to **System > Network > Interfaces** and edit the port8 interface as follows:

Alias	primary_reserved
IP/Netmask	10.11.101.101/24
Administrative Access	Ping, SSH, HTTPS, SNMP

3. Select **OK**.

You can now log into the primary unit GUI by browsing to `https://10.11.101.101`. You can also log into this primary unit CLI by using an SSH client to connect to `10.11.101.101`.

To change subordinate unit reserved management interface configuration - CLI

At this point you cannot connect to the subordinate unit reserved management interface because it does not have an IP address. Instead, this procedure describes connecting to the primary unit CLI and using the `execute ha manage` command to connect to subordinate unit CLI to change the port8 interface. You can also use a serial connection to the cluster unit CLI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. Connect to the primary unit CLI and use the `execute ha manage` command to connect to a subordinate unit CLI.

You can identify the subordinate unit from its serial number or host name. The host name appears in the CLI prompt.

2. Enter the following command to change the port8 IP address to `10.11.101.102` and set management access to HTTPS, ping, SSH, and SNMP.

```
config system interface
edit port8
set ip 10.11.101.102/24
set allowaccess https ping ssh snmp
end
```

You can now log into the subordinate unit GUI by browsing to `https://10.11.101.102`. You can also log into this subordinate unit CLI by using an SSH client to connect to `10.11.101.102`.

To configure the cluster for SNMP management using the reserved management interfaces - CLI

This procedure describes how to configure the cluster to allow the SNMP server to get status information from the primary unit and the subordinate unit. The SNMP configuration is synchronized to all cluster units. To support using the reserved management interfaces, you must add at least one HA direct management host to an SNMP community. If your SNMP configuration includes SNMP users with user names and passwords you must also enable HA direct management for SNMP users.

1. Enter the following command to add an SNMP community called `Community` and add a host to the community for the reserved management interface of each cluster unit. The host includes the IP address of the SNMP server (10.11.101.20).

```
config system snmp community
edit 1
set name Community
config hosts
edit 1
set ha-direct enable
set ip 10.11.101.20
end
end
```



Enabling ha-direct in non-HA environments makes SNMP unusable.

- 2.
3. Enter the following system command to add an SNMP user for the reserved management interface.

```
config system snmp user
edit 1
set ha-direct enable
set notify-hosts 10.11.101.20
end
```

Configure other settings as required.

To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses

From the command line of an SNMP manager, you can use the following SNMP commands to get CPU, memory and network usage information for each cluster unit. In the examples, the community name is `Community`. The commands use the MIB field names and OIDs listed below.

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.101 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.102 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Adding firewall local-in policies for the dedicated HA management interface

To add local-in policies for the dedicated management interface, enable `ha-mgmt-intf-only` and set `intf` to any. Enabling `ha-mgmt-intf-only` means the local-in policy applies only to the VDOM that contains the dedicated HA management interface. For example:

```
config firewall local-in-policy
edit 0
set ha-mgmt-intf-only enable
set intf any
set scraddr internal-net
set dstaddr mgmt-int
set action accept
set service HTTPS
set schedule weekdays
end
```

Managing individual cluster units in a virtual cluster

You can select the HA option **Do NOT Synchronize Management VDOM Configuration** if you have enabled multiple VDOMS and set a VDOM other than the root VDOM to be the management VDOM. You can select this option to prevent the management VDOM configuration from being synchronized between cluster units in a virtual cluster. This allows you to add an interface to the VDOM in each cluster unit and then to give the interfaces different IP addresses in each cluster unit, allowing you to manage each cluster unit separately.

You can also enable this feature using the following command:

```
config system ha
set standalone-mgmt-vdom enable
end
```



This feature must be disabled to manage a cluster using FortiManager.

The primary unit acts as a router for subordinate unit management traffic

HA uses routing and inter-VDOM links to route subordinate unit management traffic through the primary unit to the network. Similar to a standalone FortiGate, subordinate units may generate their own management traffic, including:

- DNS queries.
- FortiGuard Web Filtering rating requests.
- Log messages to be sent to a FortiAnalyzer unit, to a syslog server, or to the FortiGuard Analysis and Management Service.

- Log file uploads to a FortiAnalyzer unit.
- Quarantine file uploads to a FortiAnalyzer unit.
- SNMP traps.
- Communication with remote authentication servers (RADIUS, LDAP, TACACS+ and so on)

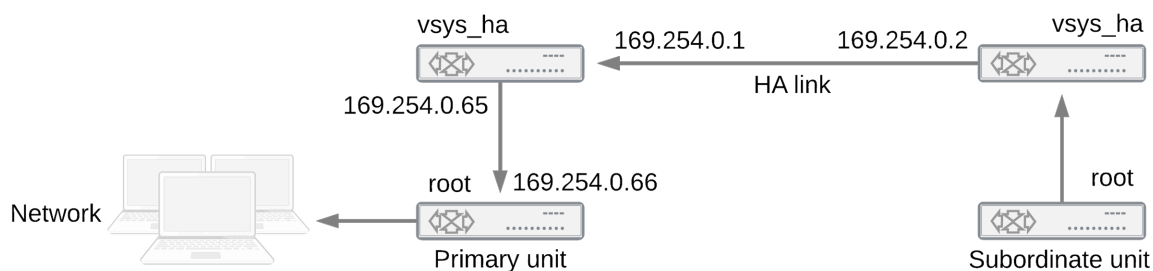
Subordinate units send this management traffic over the HA heartbeat link to the primary unit. The primary unit forwards the management traffic to its destination. The primary unit also routes replies back to the subordinate unit in the same way.

HA uses a hidden VDOM called `vsys_ha` for HA operations. The `vsys_ha` VDOM includes the HA heartbeat interfaces, and all communication over the HA heartbeat link goes through the `vsys_ha` VDOM. To provide communication from a subordinate unit to the network, HA adds hidden inter-VDOM links between the primary unit management VDOM and the primary unit `vsys_ha` VDOM. By default, `root` is the management VDOM.

Management traffic from the subordinate unit originates in the subordinate unit `vsys_ha` VDOM. The `vsys_ha` VDOM routes the management traffic over the HA heartbeat link to the primary unit `vsys_ha` VDOM. This management traffic is then routed to the primary unit management VDOM and from there out onto the network.

DNS queries and FortiGuard Web Filtering and Email Filter requests are still handled by the HA proxy so the primary unit and subordinate units share the same DNS query cache and the same FortiGuard Web Filtering and Email Filter cache. In a virtual clustering configuration, the cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering, Email Filtering, and DNS query cache.

Subordinate unit management traffic path



Cluster communication with RADIUS and LDAP servers

In an active-passive cluster, only the primary unit processes traffic, so the primary unit communicates with RADIUS or LDAP servers. In a cluster that is operating in active-active mode, subordinate units send RADIUS and LDAP requests to the primary unit over the HA heartbeat link and the primary unit routes them to their destination. The primary unit relays the responses back to the subordinate unit.

Clusters and FortiGuard services

This section describes how various FortiGate HA clustering configurations communicate with the FDN.

In an operating cluster, the primary unit communicates directly with the FortiGuard Distribution Network (FDN). Subordinate units also communicate directly with the FDN but as described above, all communication between subordinate units and the FDN is routed through the primary unit.

You must register and license all of the units in a cluster for all required FortiGuard services, both because all cluster units communicate with the FDN and because any cluster unit could potentially become the primary unit.

FortiGuard and active-passive clusters

For an active-passive cluster, only the primary unit processes traffic. Even so, all cluster units communicate with the FDN. Only the primary unit sends FortiGuard Web Filtering and Antispam requests to the FDN. All cluster units receive FortiGuard Antivirus, IPS, and application control updates from the FDN.

In an active-passive cluster the FortiGuard Web Filter and Email Filter caches are located on the primary unit in the same way as for a standalone FortiGate. The caches are not shared among cluster units so after a failover the new primary unit must build up new caches.

In an active-passive cluster all cluster units also communicate with the FortiGuard Analysis and Management Service (FAMS).

FortiGuard and active-active clusters

For an active-active cluster, both the primary unit and the subordinate units process traffic. Communication between the cluster units and the FDN is the same as for active-passive clusters with the following exception.

Because the subordinate units process traffic, they may also be making FortiGuard Web Filtering and Email Filter requests. The primary unit receives all such requests from the subordinate units and relays them to the FDN and then relays the FDN responses back to the subordinate units. The FortiGuard Web Filtering and Email Filtering URL caches are maintained on the primary unit. The primary unit caches are used for primary and subordinate unit requests.

FortiGuard and virtual clustering

For a virtual clustering configuration the management virtual domain of each cluster unit communicates with the FDN. The cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering and Email Filtering caches. All FortiGuard Web Filtering and Email Filtering requests are proxied by the management VDOM of the cluster unit that is the primary unit for the management virtual domain.

Clusters and logging

This section describes the log messages that provide information about how HA is functioning, how to view and manage logs for each unit in a cluster, and provides some example log messages that are recorded during specific cluster events.

You configure logging for a cluster in the same way as you configuring logging for a standalone FortiGate. Log configuration changes made to the cluster are synchronized to all cluster units.

All cluster units record log messages separately to the individual cluster unit's log disk, to the cluster unit's system memory, or both. You can view and manage log messages for each cluster unit from the cluster GUI Log Access page.

When remote logging is configured, all cluster units send log messages to remote FortiAnalyzer units or other remote servers as configured. HA uses routing and inter-VDOM links to route subordinate unit log traffic through the primary unit to the network.

When you configure a FortiAnalyzer unit to receive log messages from a FortiGate cluster, you should add a cluster to the FortiAnalyzer unit configuration so that the FortiAnalyzer unit can receive log messages from all cluster units.

Viewing and managing log messages for individual cluster units

This section describes how to view and manage log messages for an individual cluster unit.

To view HA cluster log messages

1. Log into the cluster GUI.
2. Go to **Log&Report > Log Config > Log Settings > GUI Preferences** and select to display logs from **Memory, Disk or FortiAnalyzer**.

For each log display, the **HA Cluster** list displays the serial number of the cluster unit for which log messages are displayed. The serial numbers are displayed in order in the list.

3. Set **HA Cluster** to the serial number of one of the cluster units to display log messages for that unit.

About HA event log messages

HA event log messages always include the host name and serial number of the cluster unit that recorded the message. HA event log messages also include the HA state of the unit and also indicate when a cluster unit switches (or moves) from one HA state to another. Cluster units can operate in the HA states listed below:

HA states

Hello	A FortiGate configured for HA operation has started up and is looking for other FortiGates with which to form a cluster.
Work	In an active-passive cluster a cluster unit is operating as the primary unit. In an active-active cluster unit is operating as the primary unit or a subordinate unit.
Standby	In an active-passive cluster the cluster unit is operating as a subordinate unit.

HA log Event log messages also indicate the virtual cluster that the cluster unit is operating in as well as the member number of the unit in the cluster. If virtual domains are not enabled, all clusters unit are always operating in virtual cluster 1. If virtual domains are enabled, a cluster unit may be operating in virtual cluster 1 or virtual cluster 2. The member number indicates the position of the cluster unit in the cluster members list. Member 0 is the primary unit. Member 1 is the first subordinate unit, member 2 is the second subordinate unit, and so on.

HA log messages

See the FortiOS log message reference for a listing of and descriptions of the HA log messages.

FortiGate HA message "HA master heartbeat interface <intf_name> lost neighbor information"

The following HA log messages may be recorded by an operating cluster:

```
2009-02-16 11:06:34 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=critical vd=root msg="HA slave heartbeat interface internal lost neighbor information"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=notice vd=root msg="Virtual cluster 1 of group 0 detected new joined HA member"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=notice vd=root msg="HA master heartbeat interface internal get peer information"
```

These log messages indicate that the cluster units could not connect to each other over the HA heartbeat link for the period of time that is given by hb-interval x hb-lost-threshold, which is 1.2 seconds with the default values.

To diagnose this problem

1. Check all heartbeat interface connections including cables and switches to make sure they are connected and operating normally.
2. Use the following commands to display the status of the heartbeat interfaces.

```
get hardware nic <heartbeat_interface_name>
diagnose hardware deviceinfo nic <heartbeat_interface_name>
```

The status information may indicate the interface status and link status and also indicate if a large number of errors have been detected.

3. If the log message only appears during peak traffic times, increase the tolerance for missed HA heartbeat packets by using the following commands to increase the lost heartbeat threshold and heartbeat interval:

```
config system ha
    set hb-lost-threshold 12
    set hb-interval 4
end
```

These settings multiply by 4 the loss detection interval. You can use higher values as well.

This condition can also occur if the cluster units are located in different buildings or even different geographical locations. Called a distributed cluster, as a result of the separation it may take a relatively long time for heartbeat packets to be transmitted between cluster units. You can support a distributed cluster by increasing the heartbeat interval so that the cluster expects extra time between heartbeat packets.

4. Optionally disable session-pickup to reduce the processing load on the heartbeat interfaces.
5. Instead of disabling session-pickup you can enable `session-pickup-delay` to reduce the number of sessions that are synchronized. With this option enabled only sessions that are active for more than 30 seconds are synchronized.

It may be useful to monitor CPU and memory usage to check for low memory and high CPU usage. You can configure event logging to monitor CPU and memory usage. You can also enable the CPU over usage and memory low SNMP events.

Once this monitoring is in place, try and determine if there have been any changes in the network or an increase of traffic recently that could be the cause. Check to see if the problem happens frequently and if so what the pattern is.

To monitor the CPU of the cluster units and troubleshoot further, use the following procedure and commands:

```
get system performance status
get system performance top 2
diagnose sys top 2
```

These commands repeated at frequent intervals will show the activity of the CPU and the number of sessions.

Search the Fortinet Knowledge Base for articles about monitoring CPU and Memory usage.

If the problem persists, gather the following information (a console connection might be necessary if connectivity is lost) and provide it to Technical Support when opening a ticket:

- Debug log from the GUI: **System > Advanced > Download Debug Log**
- CLI command output:

```
diagnose sys top 2 (keep it running for 20 seconds)
get system performance status (repeat this command multiple times to get good samples)
get system ha status
diagnose sys ha status
diagnose sys ha dump-by {all options}
diagnose netlink device list
diagnose hardware deviceinfo nic <heartbeat-interface-name>
execute log filter category 1
execute log display
```

Formatting cluster unit hard disks (log disks)

If you need to format the hard disk (also called log disk or disk storage) of one or more cluster units you should disconnect the unit from the cluster and use the `execute formatlogdisk` command to format the cluster unit hard disk then add the unit back to the cluster.

For information about how to remove a unit from a cluster and add it back, see [Disconnecting a cluster unit from a cluster on page 211](#) and [Adding a disconnected FortiGate back to its cluster on page 212](#).

Once you add the cluster unit with the formatted log disk back to the cluster you should make it the primary unit before removing other units from the cluster to format their log disks and then add them back to the cluster.

Clusters and SNMP

You can use SNMP to manage a cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration and status information and receive traps.

You configure SNMP for a cluster in the same way as configuring SNMP for a standalone FortiGate. SNMP configuration changes made to the cluster are shared by all cluster units.

Each cluster unit sends its own traps and SNMP manager systems can use SNMP get commands to query each cluster unit separately. To set SNMP get queries to each cluster unit you must create a special get command that includes the serial number of the cluster unit.

Alternatively you can use the HA reserved management interface feature to give each cluster unit a different management IP address. Then you can create an SNMP get command for each cluster unit that just includes the management IP address and does not have to include the serial number.

SNMP get command syntax for the primary unit

Normally, to get configuration and status information for a standalone FortiGate or for a primary unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to the following:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

where:

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is public.

<address_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself. The HA MIB fields and OIDs are listed below:

SNMP field names and OIDs

MIB field	OID	Description
fgHaSystemMode	.1.3.6.1.4.1.12356.101.13.1.1.0	HA mode (standalone, a-a, or a-p)
fgHaGroupId	.1.3.6.1.4.1.12356.101.13.1.2.0	The HA group ID of the cluster unit.
fgHaPriority	.1.3.6.1.4.1.12356.101.13.1.3.0	The HA priority of the cluster unit. Default 128.
fgHaOverride	.1.3.6.1.4.1.12356.101.13.1.4.0	Whether HA override is disabled or enabled for the cluster unit.
fgHaAutoSync	.1.3.6.1.4.1.12356.101.13.1.5.0	Whether automatic HA synchronization is disabled or enabled.
fgHaSchedule	.1.3.6.1.4.1.12356.101.13.1.6.0	The HA load balancing schedule. Set to none unless operating in a-p mode.
fgHaGroupName	.1.3.6.1.4.1.12356.101.13.1.7.0	The HA group name.
fgHaStatsIndex	.1.3.6.1.4.1.12356.101.13.2.1.1.1.1	The cluster index of the cluster unit. 1 for the primary unit, 2 to x for the subordinate units.
fgHaStatsSerial	.1.3.6.1.4.1.12356.101.13.2.1.1.2.1	The serial number of the cluster unit.
fgHaStatsCpuUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.3.1	The cluster unit's current CPU usage.
fgHaStatsMemUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.4.1	The cluster unit's current Memory usage.

MIB field	OID	Description
fgHaStatsNetUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.5.1	The cluster unit's current Network bandwidth usage.
fgHaStatsSesCount	.1.3.6.1.4.1.12356.101.13.2.1.1.6.1	The cluster unit's current session count.
fgHaStatsPktCount	.1.3.6.1.4.1.12356.101.13.2.1.1.7.1	The cluster unit's current packet count.
fgHaStatsByteCount	.1.3.6.1.4.1.12356.101.13.2.1.1.8.1	The cluster unit's current byte count.
fgHaStatsIdsCount	.1.3.6.1.4.1.12356.101.13.2.1.1.9.1	The number of attacks reported by the IPS for the cluster unit.
fgHaStatsAvCount	.1.3.6.1.4.1.12356.101.13.2.1.1.10.1	The number of viruses reported by the antivirus system for the cluster unit.
fgHaStatsHostname	.1.3.6.1.4.1.12356.101.13.2.1.1.11.1	The hostname of the cluster unit.

To get the HA priority for the primary unit

The following SNMP get command gets the HA priority for the primary unit. The community name is `public`. The IP address of the cluster interface configured for SNMP management access is 10.10.10.1. The HA priority MIB field is `fgHaPriority` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.1.3.0. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgHaPriority
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.1.3.0
```

SNMP get command syntax for any cluster unit

To get configuration status information for a specific cluster unit (for the primary unit or for any subordinate unit), the SNMP manager must add the serial number of the cluster unit to the SNMP get command after the community name. The community name and the serial number are separated with a dash. The syntax for this SNMP get command would be:

```
snmpget -v2c -c <community_name>-<fgt_serial> <address_ipv4> {<OID> | <MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. All units in the cluster have the same community name. The most commonly used community name is `public`.

`<fgt_serial>` is the serial number of any cluster unit. For example, FGT4002803033172. You can specify the serial number of any cluster unit, including the primary unit, to get information for that unit.

`<address_ipv4>` is the IP address of the FortiGate interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself.

If the serial number matches the serial number of a subordinate unit, the SNMP get request is sent over the HA heartbeat link to the subordinate unit. After processing the request, the subordinate unit sends the reply back over the HA heartbeat link back to the primary unit. The primary unit then forwards the response back to the SNMP manager.

If the serial number matches the serial number of the primary unit, the SNMP get request is processed by the primary unit. You can actually add a serial number to the community name of any SNMP get request. But normally you only need to do this for getting information from a subordinate unit.

To get the CPU usage for a subordinate unit

The following SNMP get command gets the CPU usage for a subordinate unit in a FortiGate-5001SX cluster. The subordinate unit has serial number FG50012205400050. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1. The HA status table MIB field is `fgHaStatsCpuUsage` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.3.1. The first command uses the MIB field name and the second uses the OID for this table:

```
snmpget -v2c -c public-FG50012205400050 10.10.10.1 fgHaStatsCpuUsage
snmpget -v2c -c public-FG50012205400050 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
```

FortiGate SNMP recognizes the community name with syntax `<community_name>-<fgt_serial>`. When the primary unit receives an SNMP get request that includes the community name followed by serial number, the FGCP extracts the serial number from the request. Then the primary unit redirects the SNMP get request to the cluster unit with that serial number. If the serial number matches the serial number of the primary unit, the SNMP get is processed by the primary unit.

Getting serial numbers of cluster units

The following SNMP get commands use the MIB field name `fgHaStatsSerial.<index>` to get the serial number of each cluster unit. Where `<index>` is the cluster unit's cluster index and 1 is the cluster index of the primary unit, 2 is the cluster index of the first subordinate unit, and 3 is the cluster index of the second subordinate unit.

The OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.2.1. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1.

The first command uses the MIB field name and the second uses the OID for this table and gets the serial number of the primary unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.1
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.1.1.2.1
```

The second command uses the MIB field name and the second uses the OID for this table and gets the serial number of the first subordinate unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.2
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.2.2
```

SNMP get command syntax - reserved management interface enabled

To get configuration and status information for any cluster unit where you have enabled the HA reserved management interface feature and assigned IP addresses to the management interface of each cluster unit, an SNMP manager would use the following get command syntax:

```
snmpget -v2c -c <community_name> <mgmt_address_ipv4> {<OID> | <MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community names to a FortiGate SNMP configuration. The most commonly used community name is `public`.

<mgmt_address_ipv4> is the IP address of the FortiGate HA reserved management interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see your FortiGate's online help.

Adding FortiClient licenses to a cluster

Each FortiGate in a cluster must have its own FortiClient license. Contact your reseller to purchase FortiClient licenses for all of the FortiGates in your cluster.

When you receive the license keys you can log into the Fortinet Support site and add the FortiClient license keys to each FortiGate. Then, as long as the cluster can connect to the Internet each cluster unit receives its FortiClient license key from the FortiGuard network.

Adding FortiClient licenses to cluster units with a reserved management interface

You can also use the following steps to manually add license keys to your cluster units from the GUI or CLI. Your cluster must be connected to the Internet and you must have configured a reserved management interface for each cluster unit.

1. Log into the GUI of each cluster unit using its reserved management interface IP address.
2. Go to the **License Information** dashboard widget and beside **FortiClient** select **Enter License**.
3. Enter the license key and select **OK**.
4. Confirm that the license has been installed and the correct number of FortiClients are licensed.
5. Repeat for all of the cluster units.

You can also use the reserved management IP address to log into each cluster unit CLI and use following command to add the license key:

```
execute FortiClient-NAC update-registration-license <license-key>
```

You can connect to the CLIs of each cluster unit using their reserved management IP address.

Adding FortiClient licenses to cluster units with no reserved management interface

If you have not set up reserved management IP addresses for your cluster units, you can still add FortiClient license keys to each cluster unit. You must log into the primary unit and then use the `execute ha manage` command to connect to each cluster unit CLI. For example, use the following steps to add a FortiClient license key a cluster of three FortiGates:

1. Log into the primary unit CLI and enter the following command to confirm the serial number of the primary unit:

```
get system status
```
2. Add the FortiClient license key for that serial number to the primary unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

You can also use the GUI to add the license key to the primary unit.
3. Enter the following command to log into the first subordinate unit:

```
execute ha manage 1
```
4. Enter the following command to confirm the serial number of the cluster unit that you have logged into:


```
get system status
```

5. Add the FortiClient license key for that serial number to the cluster unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

6. Enter the following command to log into the second subordinate unit:

```
execute ha manage 2
```

7. Enter the following command to confirm the serial number of the cluster unit that you have logged into:

```
get system status
```

8. Add the FortiClient license key for that serial number to the cluster unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

Viewing FortiClient license status and active FortiClient users for each cluster unit

To view FortiClient license status and FortiClient information for each cluster unit you must log into each cluster unit's GUI or CLI. You can do this by connecting to each cluster unit's reserved management interface if they are configured. If you have not configured reserved management interfaces you can use the `execute ha manage` command to log into each cluster unit CLI.

From the GUI, view FortiClient License status from the License Information dashboard widget and select **Details** to display the list of active FortiClient users connecting through that cluster unit. You can also see active FortiClient users by going to **User & Device > Monitor > FortiClient**.

From the CLI you can use the `execute FortiClient {list | info}` command to display FortiClient license status and active FortiClient users.

For example, use the following command to display the FortiClient license status of the cluster unit that you are logged into:

```
execute forticlient info
Maximum FortiClient connections: unlimited.
Licensed connections: 114
  NAC: 114
WANOPT: 0
Test: 0
Other connections:
  IPsec: 0
  SSLVPN: 0
```

Use the following command to display the list of active FortiClient users connecting through the cluster unit. The output shows the time the connection was established, the type of FortiClient connection, the name of the device, the user name of the person connecting, the FortiClient ID, the host operating system, and the source IP address of the session.

```
execute forticlient list
TIMESTAMP TYPE CONNECT-NAME USER CLIENT-ID HOST-OS SRC-IP
20141017 09:13:33 NAC Gordon-PC Gordon 11F76E902611484A942E31439E428C5C Microsoft
  Windows 7 , 64-bit Service Pack 1 (build 7601) 172.20.120.10
20141017 09:11:55 NAC Gordon-PC 11F76E902611484A942E31439E428C5C Microsoft Windows 7 ,
  64-bit Service Pack 1 (build 7601) 172.20.120.10
20141017 07:27:11 NAC Desktop11 Richie 9451C0B8EE3740AEB7019E920BB3761B Microsoft
  Windows 7, 64-bit Service Pack 1 (build 7601) 172.20.120.20
```

Cluster members list

To display the cluster members list, go to **System > HA**.

The cluster members list displays illustrations of the front panels of the cluster units. If the network jack for an interface is shaded green, the interface is connected. Hover the mouse pointer over each illustration to view the cluster unit host name, serial number, and how long the unit has been operating (up time). The list of monitored interfaces is also displayed.

From the cluster members list you can:

- View HA statistics.
- Use the up and down arrows to change the order in which cluster units are listed.
- See the host name of each cluster unit. To change the primary unit host name, go to the system dashboard and select Change beside the current host name in the System Information widget. To view and change a subordinate unit host name, from the cluster members list select the edit icon for a subordinate unit.
- View the status or role of each cluster unit.
- View and optionally change the HA configuration of the operating cluster.
- View and optionally change the host name and device priority of a subordinate unit.
- Disconnect a cluster unit from a cluster.
- Download the Debug log for any cluster unit. You can send this debug log file to Fortinet Technical Support to help diagnose problems with the cluster or with individual cluster units.

Virtual cluster members list

If virtual domains are enabled, you can display the cluster members list to view the status of the operating virtual clusters. The virtual cluster members list shows the status of both virtual clusters including the virtual domains added to each virtual cluster.

To display the virtual cluster members list for an operating cluster log in as the admin administrator, select Global Configuration and go to **System > HA**.

The functions of the virtual cluster members list are the same as the functions of the Cluster Members list with the following exceptions.

- When you select the edit icon for a primary unit in a virtual cluster, you can change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit and you can edit the VDOM partitioning configuration of the cluster.
- When you select the edit icon for a subordinate unit in a virtual cluster, you can change the device priority for the subordinate unit for the selected virtual cluster.

Also, the HA cluster members list changes depending on the cluster unit that you connect to.

Viewing HA statistics

From the cluster members list you can select **View HA statistics** to display the serial number, status, and monitor information for each cluster unit. To view HA statistics, go to **System > HA** and select View HA Statistics. Note the following about the HA statistics display:

- Use the serial number ID to identify each FortiGate in the cluster. The cluster ID matches the FortiGate serial number.
- Status indicates the status of each cluster unit. A green check mark indicates that the cluster unit is operating normally. A red X indicates that the cluster unit cannot communicate with the primary unit.
- The up time is the time in days, hours, minutes, and seconds since the cluster unit was last started.
- The GUI displays CPU usage for core processes only. CPU usage for management processes (for example, for HTTPS connections to the GUI) is excluded.
- The GUI displays memory usage for core processes only. Memory usage for management processes (for example, for HTTPS connections to the GUI) is excluded.

Changing the HA configuration of an operating cluster

To change the configuration settings of an operating cluster, go to **System > HA** to display the cluster members list. Select Edit for the master (or primary) unit in the cluster members list to display the HA configuration page for the cluster.

You can use the HA configuration page to check and fine tune the configuration of the cluster after the cluster is up and running. For example, if you connect or disconnect cluster interfaces you may want to change the Port Monitor configuration.

Any changes you make on this page, with the exception of changes to the device priority, are first made to the primary unit configuration and then synchronized to the subordinate units. Changing the device priority only affects the primary unit.

Changing the HA configuration of an operating virtual cluster

To change the configuration settings of the primary unit in a functioning cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to **System > HA** to display the cluster members list. Select Edit for the master (or primary) unit in virtual cluster 1 or virtual cluster 2 to display the HA configuration page for the virtual cluster.

You can use the virtual cluster HA configuration page to check and fine tune the configuration of both virtual clusters after the cluster is up and running. For example, you may want to change the Port Monitor configuration for virtual cluster 1 and virtual cluster 2 so that each virtual cluster monitors its own interfaces.

You can also use this configuration page to move virtual domains between virtual cluster 1 and virtual cluster 2. Usually you would distribute virtual domains between the two virtual clusters to balance the amount of traffic being processed by each virtual cluster.

Any changes you make on this page, with the exception of changes to the device priorities, are first made to the primary unit configuration and then synchronized to the subordinate unit.

You can also adjust device priorities to configure the role of this cluster unit in the virtual cluster. For example, to distribute traffic to both cluster units in the virtual cluster configuration, you would want one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. You can create this configuration by setting the device priorities. The cluster unit with the highest device priority in virtual cluster 1 becomes the primary unit for virtual cluster 1. The cluster unit with the highest device priority in virtual cluster 2 becomes the primary unit in virtual cluster 2.

Changing the subordinate unit host name and device priority

To change the host name and device priority of a subordinate unit in an operating cluster, go to **System > HA** to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

To change the host name and device priority of a subordinate unit in an operating cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to **System > HA** to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

You can change the host name (Peer) and device priority (Priority) of this subordinate unit. These changes only affect the configuration of the subordinate unit.

The device priority is not synchronized among cluster members. In a functioning cluster you can change device priority to change the priority of any unit in the cluster. The next time the cluster negotiates, the cluster unit with the highest device priority becomes the primary unit.

The device priority range is 0 to 255. The default device priority is 128.

Upgrading cluster firmware

You can upgrade the FortiOS firmware running on an HA cluster in the same manner as upgrading the firmware running on a standalone FortiGate. During a normal firmware upgrade, the cluster upgrades the primary unit and all subordinate units to run the new firmware image. The firmware upgrade takes place without interrupting communication through the cluster.



Upgrading cluster firmware to a new major release (for example upgrading from 5.2.6 to 5.4.1) is supported for clusters. Make sure you are taking an upgrade path described in the release notes. Even so you should back up your configuration and only perform such a firmware upgrade during a maintenance window.

To upgrade the firmware without interrupting communication through the cluster, the cluster goes through a series of steps that involve first upgrading the firmware running on the subordinate units, then making one of the subordinate units the primary unit, and finally upgrading the firmware on the former primary unit. These steps are transparent to the user and the network, but depending upon your HA configuration may result in the cluster selecting a new primary unit.

The following sequence describes in detail the steps the cluster goes through during a firmware upgrade and how different HA configuration settings may affect the outcome.

1. The administrator uploads a new firmware image from the GUI or CLI.
2. If the cluster is operating in active-active mode load balancing is turned off.
3. The cluster upgrades the firmware running on all of the subordinate units.
4. Once the subordinate units have been upgraded, a new primary unit is selected.
This primary unit will be running the new upgraded firmware.
5. The cluster now upgrades the firmware of the former primary unit.

If the age of the new primary unit is more than 300 seconds (5 minutes) greater than the age of all other cluster units, the new primary unit continues to operate as the primary unit.

This is the intended behavior but does not usually occur because the age difference of the cluster units is usually less than the cluster age difference margin of 300 seconds. So instead, the cluster negotiates again to select a primary unit as described in [Primary unit selection on page 39](#).

You can keep the cluster from negotiating again by reducing the cluster age difference margin using the `ha-uptime-diff-margin` option. However, you should be cautious when reducing the age or other problems may occur. For information about the cluster age difference margin, see [Cluster age difference margin \(grace period\) on page 42](#). For more information about changing the cluster age margin, see [Changing the cluster age difference margin on page 42](#).

6. If the cluster is operating in active-active mode, load balancing is turned back on.

Changing how the cluster processes firmware upgrades

By default cluster firmware upgrades proceed as uninterruptable upgrades that do not interrupt traffic flow. If required, you can use the following CLI command to change how the cluster handles firmware upgrades. You might want to change this setting if you are finding uninterruptable upgrades take too much time.

```
config system ha
    set uninterruptible-upgrade disable
end
```

`uninterruptible-upgrade` is enabled by default. If you disable `uninterruptible-upgrade` the cluster still upgrades the firmware on all cluster units, but all cluster units are upgraded at once; which takes less time but interrupts communication through the cluster.

Synchronizing the firmware build running on a new cluster unit

If the firmware build running on a FortiGate that you add to a cluster is older than the cluster firmware build, you may be able to use the following steps to synchronize the firmware running on the new cluster unit.

This procedure describes re-installing the same firmware build on a cluster to force the cluster to upgrade all cluster units to the same firmware build.

Due to firmware upgrade and synchronization issues, in some cases this procedure may not work. In all cases it will work to install the same firmware build on the new unit as the one that the cluster is running before adding the new unit to the cluster.

To synchronize the firmware build running on a new cluster unit

1. Obtain a firmware image that is the same as build already running on the cluster.
2. Connect to the cluster using the GUI.
3. Go to the **System Information** dashboard widget.
4. Select **Update** beside **Firmware Version**.
You can also install a newer firmware build.
5. Select **OK**.

After the firmware image is uploaded to the cluster, the primary unit upgrades all cluster units to this firmware build.

Downgrading cluster firmware

For various reasons you may need to downgrade the firmware that a cluster is running. You can use the information in this section to downgrade the firmware version running on a cluster.

In most cases you can downgrade the firmware on an operating cluster using the same steps as for a firmware upgrade. A warning message appears during the downgrade but the downgrade usually works and after the downgrade the cluster continues operating normally with the older firmware image.

Downgrading between some firmware versions, especially if features have changed between the two versions, may not always work without the requirement to fix configuration issues after the downgrade.

Only perform firmware downgrades during maintenance windows and make sure you back up your cluster configuration before the downgrade.

If the firmware downgrade that you are planning may not work without configuration loss or other problems, you can use the following downgrade procedure to make sure your configuration is not lost after the downgrade.

To downgrade cluster firmware

This example shows how to downgrade the cluster shown in Example NAT/Route mode HA network topology. The cluster consists of two cluster units (FGT_ha_1 and FGT_ha_2). The port1 and port2 interfaces are connected to networks and the port3 and port4 interfaces are connected together for the HA heartbeat.

This example, describes separating each unit from the cluster and downgrading the firmware for the standalone FortiGates. There are several ways you could disconnect units from the cluster. This example describes using the disconnect from cluster function on the cluster members list GUI page.

1. Go to the **System Information** dashboard widget and backup the cluster configuration.

From the CLI use `execute backup config`.

2. Go to **System > HA** and for FGT_ha_1 select the **Disconnect from cluster** icon.
3. Select the port2 interface and enter an IP address and netmask of 10.11.101.101/24 and select **OK**.
From the CLI you can enter the following command (FG600B3908600705 is the serial number of the cluster unit) to be able to manage the standalone FortiGate by connecting to the port2 interface with IP address and netmask 10.11.101.101/24.

```
execute ha disconnect FG600B3908600705 port2 10.11.101.101/24
```

After FGT_ha_1 is disconnected, FGT_ha_2 continues processing traffic.

4. Connect to the FGT_ha_1 GUI or CLI using IP address 10.11.101.101/24 and follow normal procedures to downgrade standalone FortiGate firmware.
5. When the downgrade is complete confirm that the configuration of FGT_ha_1 is correct.
6. Set the HA mode of FGT_ha_2 to Standalone and follow normal procedures to downgrade standalone FortiGate firmware.
Network communication will be interrupted for a short time during the downgrade.
7. When the downgrade is complete confirm that the configuration of FGT_ha_2 is correct.
8. Set the HA mode of FGT_ha_2 to Active-Passive or the required HA mode.
9. Set the HA mode of FGT_ha_1 to the same mode as FGT_ha_2.

If you have not otherwise changed the HA settings of the cluster units and if the firmware downgrades have not affected the configurations the units should negotiate and form cluster running the

downgraded firmware.

Backing up and restoring the cluster configuration

You can backup and restore the configuration of a cluster in the same way as backing up and restoring a standalone FortiGate unit. Backing up the cluster from the primary unit GUI or CLI saves a single configuration file for the cluster. If you restore this configuration file, the configuration of all cluster units is restored. The restore process keeps configuration settings of individual cluster units that are not synchronized unchanged but resets all other configuration setting to those in the restored configuration file.



When restoring the configuration of a cluster, all cluster units reboot to install the new configuration. This may result in a brief traffic interruption as all cluster units may restart at the same time.

Monitoring cluster units for failover

If the primary unit in the cluster fails, the units in the cluster renegotiate to select a new primary unit. Failure of the primary unit results in the following:

- If SNMP is enabled, the new primary unit sends HA trap messages. The messages indicate a cluster status change, HA heartbeat failure, and HA member down.
- If event logging is enabled and HA activity event is selected, the new primary unit records log messages that show that the unit has become the primary unit.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGates. The failed primary unit no longer appears on the Cluster Members list.
- The host name and serial number of the primary unit changes. You can see these changes when you log into the GUI or CLI.
- The cluster info displayed on the dashboard, cluster members list or from the `get system ha status` command changes.

If a subordinate unit fails, the cluster continues to function normally. Failure of a subordinate unit results in the following:

- If event logging is enabled and HA activity event is selected, the primary unit records log messages that show that a subordinate has been removed from the cluster.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGates. The failed unit no longer appears on the Cluster Members list.

Viewing cluster status from the CLI

Use the `get system ha status` command to display information about an HA cluster. The command displays general HA configuration settings. The command also displays information about how the cluster unit

that you have logged into is operating in the cluster.

Usually you would log into the primary unit CLI using SSH or telnet. In this case the `get system ha status` command displays information about the primary unit first, and also displays the HA state of the primary unit (the primary unit operates in the work state). However, if you log into the primary unit and then use the `execute ha manage` command to log into a subordinate unit, (or if you use a console connection to log into a subordinate unit) the `get system status` command displays information about this subordinate unit first, and also displays the HA state of this subordinate unit. The state of a subordinate unit is work for an active-active cluster and standby for an active-passive cluster.

For a virtual cluster configuration, the `get system ha status` command displays information about how the cluster unit that you have logged into is operating in virtual cluster 1 and virtual cluster 2. For example, if you connect to the cluster unit that is the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2, the output of the `get system ha status` command shows virtual cluster 1 in the work state and virtual cluster 2 in the standby state. The `get system ha status` command also displays additional information about virtual cluster 1 and virtual cluster 2.

The command includes the following fields.

Fields	Description
HA Health Status	Indicates if all cluster units are operating normally (OK) or if a problem was detected with the cluster. For example, a message similar to <code>ERROR <serial-number> is lost @ <date> <time></code> appears if one of the subordinate units leaves the cluster.
Model	The FortiGate model number.
Mode	The HA mode of the cluster, for example, HA A-P or HA A-A.
Group	The group ID of the cluster.
Debug	The debug status of the cluster.
Cluster Uptime	The number of days, hours, minutes, and seconds that the cluster has been operating.
Master selected using	Shows how the primary unit was selected the last four times that the cluster negotiated. For example, when a cluster first forms this part of the command output could have one line showing that the primary unit is the cluster unit with the highest up time. Up to four lines can be included as the cluster negotiates to choose a new primary unit on different occasions. Each line includes a time stamp and the criteria used to select the primary unit.
ses_pickup	The status of session pickup: enable or disable.
load_balance	The status of the <code>load-balance-all</code> keyword: enable or disable. Active-active clusters only.
schedule	The active-active load balancing schedule. Active-active clusters only.

Fields	Description
<code>override</code>	The status of the override option for the current cluster unit: enable or disable.
<code>Configuration Status</code>	Shows if the configurations of each of the cluster units are synchronized or not.
<code>System Usage stats</code>	Shows how busy each cluster unit is by showing the number of sessions being processed by the cluster unit, CPU usage, and memory usage.
<code>HBDEV stats</code>	Shows the status of each cluster unit's heartbeat interfaces. Includes whether the interfaces are up or down, how much data they have processed as well as errors found.
<code>Master Slave</code>	<p>Displays the host name, serial number, cluster index or priority, and role of the primary unit (master) and the subordinate units (slave).</p> <p>The order in which the cluster units are listed starts with the cluster unit that you are logged into.</p>
<code>number of vcluster</code>	The number of virtual clusters. If virtual domains are not enabled, the cluster has one virtual cluster. If virtual domains are enabled the cluster has two virtual clusters.
<code>vcluster 1</code> <code>vcluster 2</code>	The heartbeat interface IP address of the primary unit in each virtual cluster. If virtual domains are not enabled there is one vcluster and this is the IP address of the primary unit. If virtual domains are enabled then each vcluster line will have an IP address. If the IP addresses are the same then the same FortiGate is the primary unit for both virtual clusters.

Fields	Description
<code>vcluster 1</code> Master Slave	<p>The HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 1. If virtual domains are not enabled, <code>vcluster 1</code> displays information for the cluster. If virtual domains are enabled, <code>vcluster 1</code> displays information for virtual cluster 1.</p> <p>The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 1 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 1.</p> <p><code>vcluster 1</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 1. The list includes the cluster index and serial number of each cluster unit in virtual cluster 1. The cluster unit that you have logged into is at the top of the list.</p> <p>If virtual domains are not enabled and you connect to the primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the primary unit.</p> <p>If virtual domains are not enabled and you connect to a subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you have logged into.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the virtual cluster 1 primary unit.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>
<code>vcluster 2</code> Master Slave	<p><code>vcluster 2</code> only appears if virtual domains are enabled. <code>vcluster 2</code> displays the HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 2. The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 2 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 2.</p> <p><code>vcluster 2</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 2. The list includes the cluster index and serial number of each cluster unit in virtual cluster 2. The cluster unit that you have logged into is at the top of the list.</p> <p>If you connect to the virtual cluster 2 primary unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>work</code>. The display lists the cluster units starting with the virtual cluster 2 primary unit.</p> <p>If you connect to the virtual cluster 2 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>standby</code>. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>

Get system ha status example - two FortiGates in active-passive mode

The following example shows `get system ha status` output for a cluster of two FortiGate-600Ds operating in active-passive mode. The cluster is healthy and has been running for 88 days. Primary unit select took place once and the cluster has been stable since then.

The following command output was produced by connecting to the primary unit CLI (host name `External-Primary`).

```
get system ha status
HA Health Status: OK
Model: FortiGate-600D
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 88 days 07:55:15
Master selected using:
    <2016/09/20 11:45:53> FGT6HD3916800525 is selected as the master because it has the
largest value of override priority.
ses_pickup: disable
override: disable
Configuration Status:
    FGT6HD3916800525(updated 4 seconds ago): in-sync
    FGT6HD3916801195(updated 4 seconds ago): out-of-sync
System Usage stats:
    FGT6HD3916800525(updated 4 seconds ago):
        sessions=91, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=25%
    FGT6HD3916801195(updated 4 seconds ago):
        sessions=4, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=24%
HBDEV stats:
    FGT6HD3916800525(updated 4 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=7188802679/14764658/121/0, tx=14537036237/17393987/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=6599589671/9550781/122/0, tx=6599535969/9550705/0/0
    FGT6HD3916801195(updated 4 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=14537164810/17394279/118/0, tx=7188884284/14764852/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=6599649293/9550869/118/0, tx=6599632709/9550845/0/0
Master: External-Primary, FGT6HD3916800525
Slave : External-Backup , FGT6HD3916801195
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master:0 FGT6HD3916800525
Slave :1 FGT6HD3916801195
```

The following command output was produced by using `execute HA manage 0` to log into the subordinate unit CLI of the cluster shown in the previous example. The host name of the subordinate unit is `External-Primary`.

```
get system ha status
HA Health Status: OK
Model: FortiGate-600D
Mode: HA A-P
Group: 0
```

```

Debug: 0
Cluster Uptime: 88 days 08:05:42
Master selected using:
    <2016/09/20 11:45:54> FGT6HD3916800525 is selected as the master because it has the
largest value of override priority.
    <2016/09/20 11:44:23> FGT6HD3916801195 is selected as the master because it's the only
member in the cluster.
    <2016/09/20 11:44:18> FGT6HD3916801195 is selected as the master because the peer member
FGT6HD3916800525 has UPGRADE_SLAVE flag set.
    <2016/09/20 11:44:14> FGT6HD3916800525 is selected as the master because because it has
UPGRADE_MASTER flag set.
ses_pickup: disable
override: disable
Configuration Status:
    FGT6HD3916801195(updated 1 seconds ago): out-of-sync
    FGT6HD3916800525(updated 1 seconds ago): in-sync
System Usage stats:
    FGT6HD3916801195(updated 1 seconds ago):
        sessions=4, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=25%
    FGT6HD3916800525(updated 1 seconds ago):
        sessions=90, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=25%
HBDEV stats:
    FGT6HD3916801195(updated 1 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=14539719170/17400214/118/0, tx=7191448101/14770621/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=6601825943/9554019/118/0, tx=6601809359/9553995/0/0
    FGT6HD3916800525(updated 1 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=7191366421/14770426/121/0, tx=14539590448/17399920/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=6601766321/9553931/122/0, tx=6601712619/9553855/0/0
Slave : External-Backup , FGT6HD3916801195
Master: External-Primary, FGT6HD3916800525
number of vcluster: 1
vcluster 1: standby 169.254.0.2
Slave :1 FGT6HD3916801195
Master:0 FGT6HD3916800525

```

Get system ha status example - three FortiGates in active-active mode

The following example shows `get system ha status` output for a cluster of three FortiGate-5001Ds operating in active-active mode. The cluster group ID is set to 20 and session pickup is enabled. Load balance all and the load balancing schedule are set to the default value.

```

get system ha status
HA Health Status: OK
Model: FortiGate-5001D
Mode: HA A-A
Group: 20
Debug: 0
Cluster Uptime: 7 days 04:50:43
Master selected using:
    <2016/10/12 14:36:03> FG-5KD3914800284 is selected as the master because it has the
largest value of override priority.
    <2016/10/12 14:36:03> FG-5KD3914800284 is selected as the master because it has the

```

```

largest value of override priority.
<2016/10/12 13:42:46> FG-5KD3914800284 is selected as the master because it has the
largest value of override priority.
<2016/10/12 13:42:43> FG-5KD3914800353 is selected as the master because it has the
largest value of uptime.
ses_pickup: enable, ses_pickup_delay=disable
load_balance: disable
load_balance_udp: disable
schedule: Round robin.
upgrade_mode: unset
override: disable
Configuration Status:
  FG-5KD3914800284(updated 4 seconds ago): in-sync
  FG-5KD3914800353(updated 3 seconds ago): in-sync
  FG-5KD3914800344(updated 3 seconds ago): in-sync
System Usage stats:
  FG-5KD3914800284(updated 4 seconds ago):
    sessions=10, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=15%
  FG-5KD3914800353(updated 3 seconds ago):
    sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=14%
  FG-5KD3914800344(updated 3 seconds ago):
    sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=14%
HBDEV stats:
  FG-5KD3914800284(updated 4 seconds ago):
    base1: physical/1000full, up, rx-bytes/packets/dropped/errors=76249501/186982/230/0,
tx=42415292/124396/0/0
    base2: physical/1000full, up, rx-bytes/packets/dropped/errors=63858640/120488/224/0,
tx=31951552/60290/0/0
  FG-5KD3914800353(updated 3 seconds ago):
    base1: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3698904361/7099198/4636/0, tx=1750757534/3891809/0/0
    base2: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3215777622/6067509/26/0, tx=1608101090/3034153/0/0
  FG-5KD3914800344(updated 3 seconds ago):
    base1: physical/1000full, up, rx-bytes/packets/dropped/errors=58355667/132320/210/0,
tx=30477028/82533/0/0
    base2: physical/1000full, up, rx-bytes/packets/dropped/errors=50349972/95004/203/0,
tx=25126240/47408/0/0
Master: 5001d-slot4      , FG-5KD3914800284
Slave : 5001d-slot5      , FG-5KD3914800353
Slave : 5001d-slot3      , FG-5KD3914800344
number of vcluster: 1
vcluster 1: work 169.254.0.3
Master:0 FG-5KD3914800284
Slave :2 FG-5KD3914800344
Slave :1 FG-5KD3914800353

```

Get system ha status example - virtual cluster

The following example shows `get system ha status` output for a cluster of two FortiGate-5001Ds with virtual clustering enabled. The host names of the FortiGates are 5001d-slot4 and 5001d-slot5.

In this first example the `get system ha status` command was entered from 5001d_slot5. The output shows that 5001d-slot5 (serial number FG-5KD3914800353) is operating as the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2.

```
get system ha status
HA Health Status: OK
Model: FortiGate-5001D
Mode: HA A-P
Group: 20
Debug: 0
Cluster Uptime: 8 days 00:17:25
Master selected using:
  virtual cluster 1:
    <2016/10/13 09:41:42> FG-5KD3914800353 is selected as the master because it
has the largest value of serialno.
    <2016/10/12 15:12:52> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
    <2016/10/12 15:12:49> FG-5KD3914800353 is selected as the master because it
has the largest value of uptime.
    <2016/10/12 15:12:49> FG-5KD3914800353 is selected as the master because it's
the only member in the cluster.
  virtual cluster 2:
    <2016/10/13 09:48:37> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
    <2016/10/13 09:40:51> FG-5KD3914800353 is selected as the master because it
has the largest value of serialno.
ses_pickup: enable, ses_pickup_delay=disable
override: vcluster1 enable, vcluster2 enable
Configuration Status:
  FG-5KD3914800353(updated 4 seconds ago): in-sync
  FG-5KD3914800284(updated 4 seconds ago): in-sync
System Usage stats:
  FG-5KD3914800353(updated 4 seconds ago):
    sessions=10, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=14%
  FG-5KD3914800284(updated 4 seconds ago):
    sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=14%
HBDEV stats:
  FG-5KD3914800353(updated 4 seconds ago):
    base1: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3965558103/7546034/4638/0, tx=1981129551/4331175/0/0
    base2: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3432622053/6417717/27/0, tx=1823376423/3381402/0/0
  FG-5KD3914800284(updated 4 seconds ago):
    base1: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=308666532/632360/235/0, tx=307993533/571074/0/0
    base2: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=280717442/470725/229/0, tx=247222071/407529/0/0
Master: 5001d-slot5      , FG-5KD3914800353
Slave : 5001d-slot4      , FG-5KD3914800284
number of vcluster: 2
vcluster 1: work 169.254.0.1
Master:0 FG-5KD3914800353
Slave :1 FG-5KD3914800284
vcluster 2: standby 169.254.0.2
Slave :1 FG-5KD3914800353
Master:0 FG-5KD3914800284
```

The following example shows `get system ha status` output for the same cluster as shown in the previous example after using `execute ha manage 1` to log into 5001d-slot4 (serial number FG-5KD3914800284).

```
get system ha status
HA Health Status:
  ERROR: FG-5KD3914800344 is lost @ 2016/10/12 14:46:05
Model: FortiGate-5001D
Mode: HA A-P
Group: 20
Debug: 0
Cluster Uptime: 8 days 00:28:40
Master selected using:
  virtual cluster 1:
    <2016/10/13 09:41:42> FG-5KD3914800353 is selected as the master because it
has the largest value of serialno.
    <2016/10/12 15:12:52> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
    <2016/10/12 15:12:49> FG-5KD3914800353 is selected as the master because it
has the largest value of uptime.
    <2016/10/12 14:46:05> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
  virtual cluster 2:
    <2016/10/13 09:48:37> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
    <2016/10/13 09:40:51> FG-5KD3914800353 is selected as the master because it
has the largest value of serialno.
ses_pickup: enable, ses_pickup_delay=disable

override: vcluster1 enable, vcluster2 enable
Configuration Status:
  FG-5KD3914800284(updated 4 seconds ago): in-sync
  FG-5KD3914800353(updated 4 seconds ago): in-sync
System Usage stats:
  FG-5KD3914800284(updated 4 seconds ago):
    sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=14%
  FG-5KD3914800353(updated 4 seconds ago):
    sessions=10, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=14%
HBDEV stats:
  FG-5KD3914800284(updated 4 seconds ago):
    base1: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=311389146/636565/235/0, tx=310492903/575295/0/0
    base2: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=283083317/474100/229/0, tx=249587946/410904/0/0
  FG-5KD3914800353(updated 4 seconds ago):
    base1: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3968057548/7550256/4638/0, tx=1983852314/4335382/0/0
    base2: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=3434987928/6421092/27/0, tx=1825742298/3384777/0/0
Slave : 5001d-slot4      , FG-5KD3914800284
Master: 5001d-slot5     , FG-5KD3914800353
number of vcluster: 2
vcluster 1: standby 169.254.0.1
Slave :1 FG-5KD3914800284
```

```
Master:0 FG-5KD3914800353
vcluster 2: work 169.254.0.2
Master:0 FG-5KD3914800284
Slave :1 FG-5KD3914800353
```

About the HA cluster index and the execute ha manage command

When a cluster starts up, the FortiGate Cluster Protocol (FGCP) assigns a cluster index and a HA heartbeat IP address to each cluster unit based on the serial number of the cluster unit. The FGCP selects the cluster unit with the highest serial number to become the primary unit. The FGCP assigns a cluster index of 0 and an HA heartbeat IP address of 169.254.0.1 to this unit. The FGCP assigns a cluster index of 1 and an HA heartbeat IP address of 169.254.0.2 to the cluster unit with the second highest serial number. If the cluster contains more units, the cluster unit with the third highest serial number is assigned a cluster index of 2 and an HA heartbeat IP address of 169.254.0.3, and so on. You can display the cluster index assigned to each cluster unit using the `get system ha status` command. Also when you use the `execute ha manage` command you select a cluster unit to log into by entering its cluster index.

The cluster index and HA heartbeat IP address only change if a unit leaves the cluster or if a new unit joins the cluster. When one of these events happens, the FGCP resets the cluster index and HA heartbeat IP address of each cluster unit according to serial number in the same way as when the cluster first starts up.

Each cluster unit keeps its assigned cluster index and HA heartbeat IP address even as the units take on different roles in the cluster. After the initial cluster index and HA heartbeat IP addresses are set according to serial number, the FGCP checks other primary unit selection criteria such as device priority and monitored interfaces. Checking these criteria could result in selecting a cluster unit without the highest serial number to operate as the primary unit.

Even if the cluster unit without the highest serial number now becomes the primary unit, the cluster indexes and HA heartbeat IP addresses assigned to the individual cluster units do not change. Instead the FGCP assigns a second cluster index, which could be called the operating cluster index, to reflect this role change. The operating cluster index is 0 for the primary unit and 1 and higher for the other units in the cluster. By default both sets of cluster indexes are the same. But if primary unit selection selects the cluster unit that does not have the highest serial number to be the primary unit then this cluster unit is assigned an operating cluster index of 0. The operating cluster index is used by the FGCP only. You can display the operating cluster index assigned to each cluster unit using the `get system ha status` command. There are no CLI commands that reference the operating cluster index.



Even though there are two cluster indexes there is only one HA heartbeat IP address and the HA heartbeat address is not affected by a change in the operating cluster index.

Using the execute ha manage command

When you use the CLI command `execute ha manage <index_integer>` to connect to the CLI of another cluster unit, the `<index_integer>` that you enter is the cluster index of the unit that you want to connect to.

Using get system ha status to display cluster indexes

You can display the cluster index assigned to each cluster unit using the CLI command `get system ha status`. The following example shows the information displayed by the `get system ha status` command

for a cluster consisting of two FortiGates operating in active-passive HA mode with virtual domains not enabled and without virtual clustering.

```
get system ha status
.
.
.
Slave :1 FGT6HD3916801195
Master:0 FGT6HD3916800525
```

In this example, the cluster unit with serial number FG50012205400050 has the highest serial number and so has a cluster index of 0 and the cluster unit with serial number FG50012204400045 has a cluster index of 1. From the CLI of the primary (or master) unit of this cluster you can connect to the CLI of the subordinate (or slave) unit using the following command:

```
execute ha manage 1
```

This works because the cluster unit with serial number FG50012204400045 has a cluster index of 1.

The `get system ha status` command output shows two similar lists of indexes and serial numbers. The listing on the sixth and seventh lines of the command output are the cluster indexes assigned according to cluster unit serial number. These are the cluster indexes that you enter when using the `execute ha manage` command. The cluster indexes shown in the last two lines of the command output are the operating cluster indexes that reflect how the cluster units are actually operating in the cluster. In this example both sets of cluster indexes are the same.

The last three lines of the command output display the status of vcluster 1. In a cluster consisting of two cluster units operating without virtual domains enabled all clustering actually takes place in virtual cluster 1. HA is designed to work this way to support virtual clustering. If this cluster was operating with virtual domains enabled, adding virtual cluster 2 is similar to adding a new copy of virtual cluster 1. Virtual cluster 2 is visible in the `get system ha status` command output when you add virtual domains to virtual cluster 2.

The HA heartbeat IP address displayed on line 8 is the HA heartbeat IP address of the cluster unit that is actually operating as the primary unit. For a default configuration this IP address will always be 169.254.0.1 because the cluster unit with the highest serial number will be the primary unit. This IP address changes if the operating primary unit is not the primary unit with the highest serial number.

Example actual and operating cluster indexes do not match

This example shows `get system ha status` command output for same cluster of two FortiGate-5001SX units. However, in this example the device priority of the cluster unit with the serial number FG50012204400045 is increased to 200. As a result the cluster unit with the lowest serial number becomes the primary unit. This means the actual and operating cluster indexes of the cluster units do not match.

```
get system ha status
.
.
.
Master:1 FG50012205400050
Slave :0 FG50012204400045
```

The actual cluster indexes have not changed but the operating cluster indexes have. Also, the HA heartbeat IP address displayed for vcluster 1 has changed to 169.254.0.2.

Virtual clustering example output

The `get system ha status` command output is the same if a cluster is operating with virtual clustering turned on but with all virtual domains in virtual cluster 1. The following `get system ha status` command output example shows the same cluster operating as a virtual cluster with virtual domains in virtual cluster 1 and added to virtual cluster 2. In this example the cluster unit with serial number FG50012204400045 is the primary unit for virtual cluster 1 and the cluster unit with serial number FG50012205400050 is the primary unit for virtual cluster 2.

```
get system ha status
.
.
.
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:1 FG50012205400050
Slave :0 FG50012204400045
vcluster 2: standby 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045
```

This example shows three sets of indexes. The indexes in lines six and seven are still used by the `execute ha manage` command. The indexes on lines ten and eleven are for the primary and subordinate units in virtual cluster 1 and the indexes on the last two lines are for virtual cluster 2.

Managing individual cluster units

The following procedure describes how to use SSH to log into the primary unit CLI and from there to use the `execute ha manage` command to connect to the CLI of any other unit in the cluster. The procedure is very similar if you use telnet, or the GUI dashboard CLI console.

You can use the `execute ha manage` command from the CLI of any cluster unit to log into the CLI of another the cluster unit. Usually you would use this command from the CLI of the primary unit to log into the CLI of a subordinate unit. However, if you have logged into a subordinate unit CLI, you can use this command to log into the primary unit CLI, or the CLI of another subordinate unit.

Using SSH or telnet or the GUI dashboard CLI console you can only log into the primary unit CLI. Using a direct console connection you can log into any cluster unit. In both cases you can use `execute ha manage` to connect to the CLI of other cluster units.

1. Use SSH to connect to the cluster and log into the primary unit CLI.
Connect to any cluster interface configured for SSH administrative access to log into the cluster.
2. Enter the following command followed by a space and type a question mark (?):
`execute ha manage`
The CLI displays a list of all the subordinate units in the cluster. Each cluster unit is numbered, starting at 1. The information displayed for each cluster unit includes the unit serial number and the host name of the unit.
3. Complete the command with the number of the subordinate unit to log into. For example, to log into subordinate unit 1, enter the following command:
`execute ha manage 1`
Press Enter to connect to and use an administrator account to log into the CLI of the selected

subordinate unit. If this subordinate unit has a different host name, the CLI prompt changes to this host name.

You can use CLI commands to manage this subordinate unit. If you make changes to the configuration of any cluster unit (primary or subordinate unit) these changes are synchronized to all cluster units.

4. You can now use the `execute ha manage` command to connect to any other cluster unit (including the primary unit). You can also use the `exit` command to return to the primary unit CLI.

Disconnecting a cluster unit from a cluster

Use the following procedures to disconnect a cluster unit from a functioning cluster without disrupting the operation of the cluster. You can disconnect a cluster unit if you need to use the disconnected FortiGate for another purpose, such as to act as a standalone firewall.

You can use the following procedures for a standard cluster and for a virtual clustering configuration. To use the following procedures from a virtual cluster you must be logged in as the admin administrator and you must have selected Global Configuration.

When you disconnect a cluster unit you must assign an IP address and netmask to one of the interfaces of the disconnected unit. You can disconnect any unit from the cluster even the primary unit. After the unit is disconnected, the cluster responds as if the disconnected unit has failed. The cluster may renegotiate and may select a new primary unit.

When the cluster unit is disconnected the HA mode is changed to standalone. In addition, all interface IP addresses of the disconnected unit are set to 0.0.0.0 except for the interface that you configure.

Otherwise the configuration of the disconnected unit is not changed. The HA configuration of the disconnected unit is not changed either (except to change the HA mode to Standalone).

To disconnect a cluster unit from a cluster - GUI

1. Go to **System > HA** to view the cluster members list.
2. Select the Disconnect from cluster icon for the cluster unit to disconnect from the cluster.
3. Select the interface that you want to configure. You also specify the IP address and netmask for this interface. When the FortiGate is disconnected, all management access options are enabled for this interface.
4. Specify an IP address and netmask for the interface. You can use this IP address to connect to the interface to configure the disconnected FortiGate.
5. Select **OK**.

The FortiGate is disconnected from the cluster and the cluster may renegotiate and select a new primary unit. The selected interface of the disconnected unit is configured with the specified IP address and netmask.

To disconnect a cluster unit from a cluster - CLI

1. Enter the following command to disconnect a cluster unit with serial number FGT5002803033050. The internal interface of the disconnected unit is set to IP address 1.1.1.1 and netmask 255.255.255.0.

```
execute ha disconnect FGT5002803033050 internal 1.1.1.1 255.255.255.0
```

Adding a disconnected FortiGate back to its cluster

If you disconnect a FortiGate from a cluster, you can re-connect the disconnected FortiGate to the cluster by setting the HA mode of the disconnected unit to match the HA mode of the cluster. Usually the disconnected unit rejoins the cluster as a subordinate unit and the cluster automatically synchronizes its configuration.



You do not have to change the HA password on the disconnected unit unless the HA password has been changed after the unit was disconnected. Disconnecting a unit from a cluster does not change the HA password.



You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. You should also make sure that the HA `override` CLI option is not enabled on the disconnected unit. Otherwise, when the disconnected unit joins the cluster, the cluster will renegotiate and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units. This configuration change might disrupt the operation of the cluster.

The following procedure assumes that the disconnected FortiGate is correctly physically connected to your network and to the cluster but is not running in HA mode and not part of the cluster.

Before you start this procedure you should note the device priority of the primary unit.

To add a disconnected FortiGate back to its cluster - GUI

1. Log into the disconnected FortiGate.
If virtual domains are enabled, log in as the admin administrator and select Global Configuration.
2. Go to **System > HA**.
3. Change Mode to match the mode of the cluster.
4. If required, change the group name and password to match the cluster.
5. Set the Device Priority lower than the device priority of the primary unit.
6. Select **OK**.
The disconnected FortiGate joins the cluster.

To add a disconnected FortiGate back to its cluster - CLI

1. Log into the CLI of the FortiGate to be added back to the cluster.
2. Enter the following command to access the global configuration and add the FortiGate back to a cluster operating in active-passive mode and set the device priority to 50 (a low number) so that this unit will not become the primary unit:

```
config global
  config system ha
    set mode a-p
    set priority 50
  end
end
```

You may have to also change the group name, group id and password. However if you have not

changed these for the cluster or the FortiGate after it was disconnected from the cluster you should not have to adjust them now.

HA diagnose commands

You can use the following diagnose command to display a data about a cluster:

```
diagnose sys ha dump-by {all-xdb | all-vcluster| rcache | all-group |
memory | debug-zone | vdom | kernel | device | stat| sesync}
```

The example out put below is from a cluster of two FortiGate-5001Cs. In this cluster the base1 and base2 interfaces communicate the HA heartbeat and port monitoring has been added to port1.

all-xdb

This command displays information about the current configuration of the cluster and how its operating. You can use the out to determine the primary unit, the state of port monitoring as well as most cluster configuration details and status.

```
diagnose sys ha dump-by all-xdb
HA information.
idx=1,nxentry=2,linkfails=7,flags=0,digest=7.72.e3.2e.8e.d1...
xentry FG-5KC3E13800046 nhbdev=2,nventry=0, hops=0.
base1, 50, mac=0.9.f,bc.e.6c, neighbor=1.
id=FG-5KC3E13800084, mac=0.9.f,bc.11.18.
base2, 50, mac=0.9.f,bc.e.71, neighbor=1.
id=FG-5KC3E13800084, mac=0.9.f,bc.11.1d.

xentry FG-5KC3E13800084 nhbdev=2,nventry=1, hops=1.
base1, 50, mac=0.9.f,bc.11.18, neighbor=1.
id=FG-5KC3E13800046, mac=0.9.f,bc.e.6c.
base2, 50, mac=0.9.f,bc.11.1d, neighbor=1.
id=FG-5KC3E13800046, mac=0.9.f,bc.e.71.
npath=1,FG-5KC3E13800084
ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,flag=0,time=12974,mon=0
mondev=port1,50

idx=0,nxentry=2,linkfails=7,flags=3,digest=7.95.b.9.a8.5d...
xentry FG-5KC3E13800084 nhbdev=2,nventry=1, hops=0.
base1, 50, mac=0.9.f,bc.11.18, neighbor=1.
id=FG-5KC3E13800046, mac=0.9.f,bc.e.6c.
base2, 50, mac=0.9.f,bc.11.1d, neighbor=1.
id=FG-5KC3E13800046, mac=0.9.f,bc.e.71.
ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,flag=0,time=12974,mon=0
mondev=port1,50

xentry FG-5KC3E13800046 nhbdev=2,nventry=1, hops=1.
base1, 50, mac=0.9.f,bc.e.6c, neighbor=1.
id=FG-5KC3E13800084, mac=0.9.f,bc.11.18.
base2, 50, mac=0.9.f,bc.e.71, neighbor=1.
id=FG-5KC3E13800084, mac=0.9.f,bc.11.1d.
npath=1,FG-5KC3E13800046
ventry idx=0,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,override=0,flag=0,time=2,mon=0
mondev=port1,50
```

all-vcluster

This command displays the status and configuration of the individual cluster units. You can use the output of this command to determine the primary unit and the status of each cluster unit.

```
diagnose sys ha dump-by all-vcluster
HA information.
vcluster id=1, nventry=2, state=work, digest=5.f8.d1.63.4d.d2...
ventry idx=0,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,override=0,flag=1,time=0,mon=0
mondev=port1,50
ventry idx=1,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,flag=0,time=12974,mon=0
```

stat

This command displays some statistics about how well the cluster is functioning. Information includes packet counts, memory use, failed links and ping failures.

```
diagnose sys ha dump-by stat
HA information.
packet count = 1, memory = 220.
check_linkfails = 0, linkfails = 0, check_pingsvrfails = 2822
bufcnt = -5, bufmem = 0
```

HA and failover protection

In FortiGate active-passive HA, the FortiGate Clustering Protocol (FGCP) provides failover protection. This means that an active-passive cluster can provide FortiGate services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiGate. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in a mission-critical environment.

The FGCP supports three kinds of failover protection. Device failover automatically replaces a failed device and restarts traffic flow with minimal impact on the network. Link failover maintains traffic flow if a link fails. Session failover resumes communication sessions with minimal loss of data if a device or link failover occurs.

This chapter describes how FGCP failover protection works and provides detailed NAT/Route and Transparent mode packet flow descriptions.

About active-passive failover

To achieve failover protection in an active-passive cluster, one of the cluster units functions as the primary unit, while the rest of the cluster units are subordinate units, operating in an active stand-by mode. The cluster IP addresses and HA virtual MAC addresses are associated with the cluster interfaces of the primary unit. All traffic directed at the cluster is actually sent to and processed by the primary unit.

While the cluster is functioning, the primary unit functions as the FortiGate network security device for the networks that it is connected to. In addition, the primary unit and subordinate units use the HA heartbeat to keep in constant communication. The subordinate units report their status to the cluster unit and receive and store connection and state table updates.

Device failure

If the primary unit encounters a problem that is severe enough to cause it to fail, the remaining cluster units negotiate to select a new primary unit. This occurs because all of the subordinate units are constantly waiting to negotiate to become primary units. Only the heartbeat packets sent by the primary unit keep the subordinate units from becoming primary units. Each received heartbeat packet resets negotiation timers in the subordinate units. If this timer is allowed to run out because the subordinate units do not receive heartbeat packets from the primary unit, the subordinate units assume that the primary unit has failed, and negotiate to become primary units themselves.

Using the same FGCP negotiation process that occurs when the cluster starts up, after they determine that the primary unit has failed, the subordinate units negotiate amongst themselves to select a new primary unit. The subordinate unit that wins the negotiation becomes the new primary unit with the same MAC and IP addresses as the former primary unit. The new primary unit then sends gratuitous ARP packets out all of its interfaces to inform attached switches to send traffic to the new primary unit. Sessions then resume with the new primary unit.

Link failure

If a primary unit interface fails or is disconnected while a cluster is operation, a link failure occurs. When a link failure occurs the cluster units negotiate to select a new primary unit. Since the primary unit has not stopped

operating, it participates in the negotiation. The link failure means that a new primary unit must be selected and the cluster unit with the link failure joins the cluster as a subordinate unit.

Just as for a device failover, the new primary unit sends gratuitous arp packets out all of its interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary unit.

If a subordinate unit experiences a device failure its status in the cluster does not change. However, in future negotiations a cluster unit with a link failure is unlikely to become the primary unit.

Primary unit recovery

If a primary unit recovers after a device or link failure, it will operate as a subordinate unit, unless the `override` CLI keyword is enabled and its device priority is set higher than the unit priority of other cluster units (see [HA override on page 47](#)).

About active-active failover

HA failover in a cluster running in active-active mode is similar to active-passive failover described above. Active-active subordinate units are constantly waiting to negotiate to become primary units and, if session failover is enabled, continuously receive connection state information from the primary unit. If the primary unit fails, or one of the primary unit interfaces fails, the cluster units use the same mechanisms to detect the failure and to negotiate to select a new primary unit. If session failover is enabled, the new primary unit also maintains communication sessions through the cluster using the shared connection state table.

Active-active HA load balances sessions among all cluster units. For session failover, the cluster must maintain all of these sessions. To load balance sessions, the functioning cluster uses a load balancing schedule to distribute sessions to all cluster units. The shared connection state table tracks the communication sessions being processed by all cluster units (not just the primary unit). After a failover, the new primary unit uses the load balancing schedule to re-distribute all of the communication sessions recorded in the shared connection state table among all of the remaining cluster units. The connections continue to be processed by the cluster, but possibly by a different cluster unit, and are handled according to their last known state.

Device failover

The FGCP provides transparent device failover. Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

In the case of FortiOS HA, the device is the primary unit. If the primary unit fails, device failover ensures that one of the subordinate units in the cluster automatically takes the place of the primary unit and can continue processing network traffic in the same way as the failed primary unit.



Device failover does not maintain communication sessions. After a device failover, communication sessions have to be restarted. To maintain communication sessions, you must enable session failover. See [Device failover on page 216](#).

FortiGate HA device failover is supported by the HA heartbeat, virtual MAC addresses, configuration synchronization, route synchronization and IPsec VPN SA synchronization.

The HA heartbeat makes sure that the subordinate units detect a primary unit failure. If the primary unit fails to respond on time to HA heartbeat packets the subordinate units assume that the primary unit has failed and negotiate to select a new primary unit.

The new primary unit takes the place of the failed primary unit and continues functioning in the same way as the failed primary unit. For the new primary unit to continue functioning like the failed primary unit, the new primary unit must be able to reconnect to network devices and the new primary unit must have the same configuration as the failed primary unit.

FortiGate HA uses virtual MAC addresses to reconnect the new primary unit to network devices. The FGCP causes the new primary unit interfaces to acquire the same virtual MAC addresses as the failed primary unit. As a result, the new primary unit has the same network identity as the failed primary unit.

The new primary unit interfaces have different physical connections than the failed primary unit. Both the failed and the new primary unit interfaces are connected to the same switches, but the new primary unit interfaces are connected to different ports on these switches. To make sure that the switches send packets to the new primary unit, the new primary unit interfaces send gratuitous ARP packets to the connected switches. These gratuitous ARP packets notify the switches that the primary unit MAC and IP addresses are on different switch ports and cause the switches to send packets to the ports connected to the new primary unit. In this way, the new primary unit continues to receive packets that would otherwise have been sent to the failed primary unit.

Configuration synchronization means that the new primary unit always has the same configuration as the failed primary unit. As a result the new primary unit operates in exactly the same way as the failed primary unit. If configuration synchronization were not available the new primary unit may not process network traffic in the same way as the failed primary unit.

Kernel routing table synchronization synchronizes the primary unit kernel routing table to all subordinate units so that after a failover the new primary unit does not have to form a completely new routing table. IPsec VPN SA synchronization synchronizes IPsec VPN security associations (SAs) and other IPsec session data so that after a failover the new primary unit can resume IPsec tunnels without having to establish new SAs.

HA heartbeat and communication between cluster units

The HA heartbeat keeps cluster units communicating with each other. The heartbeat consists of hello packets that are sent at regular intervals by the heartbeat interface of all cluster units. These hello packets describe the state of the cluster unit and are used by other cluster units to keep all cluster units synchronized.

HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The default time interval between HA heartbeats is 200 ms. The FGCP uses link-local IPv4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

For best results, isolate the heartbeat devices from your user networks by connecting the heartbeat devices to a separate switch that is not connected to any network. If the cluster consists of two FortiGates you can connect the heartbeat device interfaces directly using a crossover cable. Heartbeat packets contain sensitive information about the cluster configuration. Heartbeat packets may also use a considerable amount of network bandwidth. For these reasons, it is preferable to isolate heartbeat packets from your user networks.

On startup, a FortiGate configured for HA operation broadcasts HA heartbeat hello packets from its HA heartbeat interface to find other FortiGates configured to operate in HA mode. If two or more FortiGates operating in HA

mode connect with each other, they compare HA configurations (HA mode, HA password, and HA group ID). If the HA configurations match, the units negotiate to form a cluster.

While the cluster is operating, the HA heartbeat confirms that all cluster units are functioning normally. The heartbeat also reports the state of all cluster units, including the communication sessions that they are processing.

Heartbeat interfaces

A heartbeat interface is an Ethernet network interface in a cluster that is used by the FGCP for HA heartbeat communications between cluster units.

To change the HA heartbeat configuration go to **System > HA** and select the *FortiGate interfaces* to use as HA heartbeat interfaces.



Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.

From the CLI enter the following command to make port4 and port5 HA heartbeat interfaces and give both interfaces a heartbeat priority of 150:

```
config system ha
  set hbdev port4 150 port5 150
end
```

The following example shows how to change the default heartbeat interface configuration so that the port4 and port1 interfaces can be used for HA heartbeat communication and to give the port4 interface the highest heartbeat priority so that port4 is the preferred HA heartbeat interface.

```
config system ha
  set hbdev port4 100 port1 50
end
```

By default, for most FortiGate models two interfaces are configured to be heartbeat interfaces. You can change the heartbeat interface configuration as required. For example you can select additional or different heartbeat interfaces. You can also select only one heartbeat interface.

In addition to selecting the heartbeat interfaces, you also set the **Priority** for each heartbeat interface. In all cases, the heartbeat interface with the highest priority is used for all HA heartbeat communication. If the interface fails or becomes disconnected, the selected heartbeat interface that has the next highest priority handles all heartbeat communication.

If more than one heartbeat interface has the same priority, the heartbeat interface with the highest priority that is also highest in the heartbeat interface list is used for all HA heartbeat communication. If this interface fails or becomes disconnected, the selected heartbeat interface with the highest priority that is next highest in the list handles all heartbeat communication.

The default heartbeat interface configuration sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration if one or both of the default heartbeat interfaces are connected. You can select different heartbeat interfaces, select more heartbeat interfaces and change heartbeat priorities according to your requirements.

For the HA cluster to function correctly, you must select at least one heartbeat interface and this interface of all of the cluster units must be connected together. If heartbeat communication is interrupted and cannot failover to a second heartbeat interface, the cluster units will not be able to communicate with each other and more than one

cluster unit may become a primary unit. As a result the cluster stops functioning normally because multiple devices on the network may be operating as primary units with the same IP and MAC addresses creating a kind of split brain scenario.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0. The higher the number the higher the priority.

In most cases you can maintain the default heartbeat interface configuration as long as you can connect the heartbeat interfaces together. Configuring HA heartbeat interfaces is the same for virtual clustering and for standard HA clustering.

You can enable heartbeat communications for physical interfaces, but not for VLAN subinterfaces, IPsec VPN interfaces, redundant interfaces, or for 802.3ad aggregate interfaces. You cannot select these types of interfaces in the heartbeat interface list.

Selecting more heartbeat interfaces increases reliability. If a heartbeat interface fails or is disconnected, the HA heartbeat fails over to the next heartbeat interface.

You can select up to 8 heartbeat interfaces. This limit only applies to FortiGates with more than 8 physical interfaces.

HA heartbeat traffic can use a considerable amount of network bandwidth. If possible, enable HA heartbeat traffic on interfaces used only for HA heartbeat traffic or on interfaces connected to less busy networks.

Connecting HA heartbeat interfaces

For most FortiGate models if you do not change the heartbeat interface configuration, you can isolate the default heartbeat interfaces of all of the cluster units by connecting them all to the same switch. Use one switch per heartbeat interface. If the cluster consists of two units you can connect the heartbeat interfaces together using crossover cables.

HA heartbeat and data traffic are supported on the same cluster interface. In NAT/Route mode, if you decide to use heartbeat interfaces for processing network traffic or for a management connection, you can assign the interface any IP address. This IP address does not affect HA heartbeat traffic.

In Transparent mode, you can connect the heartbeat interface to your network and enable management access. You would then establish a management connection to the interface using the Transparent mode management IP address. This configuration does not affect HA heartbeat traffic.

Heartbeat packets and heartbeat interface selection

HA heartbeat hello packets are constantly sent by all of the enabled heartbeat interfaces. Using these hello packets, each cluster unit confirms that the other cluster units are still operating. The FGCP selects one of the heartbeat interfaces to be used for communication between the cluster units. The FGCP selects the heartbeat interface for heartbeat communication based on the linkfail states of the heartbeat interfaces, on the priority of the heartbeat interfaces, and on the interface index.

The FGCP checks the linkfail state of all heartbeat interfaces to determine which ones are connected. The FGCP selects one of these connected heartbeat interfaces to be the one used for heartbeat communication. The FGCP selects the connected heartbeat interface with the highest priority for heartbeat communication.

If more than one connected heartbeat interface has the highest priority the FGCP selects the heartbeat interface with the lowest interface index. The GUI lists the FortiGate interfaces in alphabetical order. This order corresponds to the interface index order with lowest index at the top and highest at the bottom. If more than one

heartbeat interface has the highest priority, the FGCP selects the interface that is highest in the heartbeat interface list (or first in alphabetical order) for heartbeat communication.

If the interface that is processing heartbeat traffic fails or becomes disconnected, the FGCP uses the same criteria to select another heartbeat interface for heartbeat communication. If the original heartbeat interface is fixed or reconnected, the FGCP again selects this interface for heartbeat communication.

The HA heartbeat communicates cluster session information, synchronizes the cluster configuration, synchronizes the cluster kernel routing table, and reports individual cluster member status. The HA heartbeat constantly communicates HA status information to make sure that the cluster is operating properly.

Interface index and display order

The GUI and CLI display interface names in alphanumeric order. For example, the sort order for a FortiGate with 10 interfaces (named port1 through port10) places port10 at the bottom of the list:

- port1
- port2 through 9
- port10

However, interfaces are indexed in hash map order, rather than purely by alphabetic order or purely by interface number value comparisons. As a result, the list is sorted primarily alphabetical by interface name (for example, base1 is before port1), then secondarily by index numbers:

- port1
- port10
- port2 through port9

HA heartbeat interface IP addresses

The FGCP uses link-local IPv4 addresses ([RFC 3927](#)) in the 169.254.0.x range for HA heartbeat interface IP addresses and for inter-VDOM link interface IP addresses. When a cluster initially starts up, the primary unit heartbeat interface IP address is 169.254.0.1. Subordinate units are assigned heartbeat interface IP addresses in the range 169.254.0.2 to 169.254.0.63. HA inter-VDOM link interfaces on the primary unit are assigned IP addresses 169.254.0.65 and 169.254.0.66.

If a failover occurs, the primary unit heartbeat interface could be something other than 169.254.0.1. If for example, the first subordinate unit is now the primary unit, the primary unit heartbeat interface IP address would be 169.254.0.2.

The output from the `get system ha status` CLI command shows the HA heartbeat interface IP address of the primary unit.

```
get system ha status
.
.
.
vcluster 1: work 169.254.0.2
.
.
.
```

You can also use the `execute traceroute` command from the subordinate unit CLI to display HA heartbeat IP addresses and the HA inter-VDOM link IP addresses. For example, use `execute ha manage 1` to connect to the subordinate unit CLI and then enter the following command to trace the route to an IP address on your network:

```
execute traceroute 172.20.20.10
traceroute to 172.20.20.10 (172.20.20.10), 32 hops max, 72 byte packets
 1 169.254.0.1 0 ms 0 ms 0 ms
 2 169.254.0.66 0 ms 0 ms 0 ms
 3 172.20.20.10 0 ms 0 ms 0 ms
```

Both HA heartbeat and data traffic are supported on the same FortiGate interface. All heartbeat communication takes place on a separate VDOM called `vsys_ha`. Heartbeat traffic uses a virtual interface called `port_ha` in the `vsys_ha` VDOM. Data and heartbeat traffic use the same physical interface, but they're logically separated into separate VDOMs.

Heartbeat packet Ethertypes

Normal IP packets are 802.3 packets that have an Ethernet type (Ethertype) field value of 0x0800. Ether type values other than 0x0800 are understood as level 2 frames rather than IP packets.

By default, HA heartbeat packets use the following Ethertypes:

- HA heartbeat packets for NAT/Route mode clusters use Ether type 0x8890. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `ha-eth-type` option of the `config system ha` command.
- HA heartbeat packets for Transparent mode clusters use Ether type 0x8891. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `hc-eth-type` option of the `config system ha` command.
- HA telnet sessions between cluster units over HA heartbeat links use Ether type 0x8893. The telnet sessions are used to synchronize the cluster configurations. Telnet sessions are also used when an administrator uses the `execute ha manage` command to connect from one cluster unit CLI to another. You can change the Ether type of these packets using the `l2ep-eth-type` option of the `config system ha` command.

Because heartbeat packets are recognized as level 2 frames, the switches and routers on your heartbeat network that connect to heartbeat interfaces must be configured to allow them. If level2 frames are dropped by these network devices, heartbeat traffic will not be allowed between the cluster units.

Some third-party network equipment may use packets with these Ethertypes for other purposes. For example, Cisco N5K/Nexus switches use Ether type 0x8890 for some functions. When one of these switches receives Ether type 0x8890 packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGates connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

Alternatively, you can use the following CLI options to change the Ethertypes of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For example, use the following command to change the Ether type of the HA heartbeat packets from 0x8890 to 0x8895 and to change the Ether type of HA Telnet session packets from 0x8891 to 0x889f:

```
config system ha
  set ha-eth-type 8895
  set l2ep-eth-type 889f
end
```

Modifying heartbeat timing

In an HA cluster, if a cluster unit CPU becomes very busy, the cluster unit may not be able to send heartbeat packets on time. If heartbeat packets are not sent on time other units in the cluster may think that the cluster unit has failed and the cluster will experience a failover.

A cluster unit CPU may become very busy if the cluster is subject to a syn flood attack, if network traffic is very heavy, or for other similar reasons. You can use the following CLI commands to configure how the cluster times HA heartbeat packets:

```
config system ha
  set hb-interval <interval_integer>
  set hb-lost-threshold <threshold_integer>
  set hello-holddown <holddown_integer>
end
```

Changing the lost heartbeat threshold

The lost heartbeat threshold is the number of consecutive heartbeat packets that are not received from another cluster unit before assuming that the cluster unit has failed. The default value is 6, meaning that if the 6 heartbeat packets are not received from a cluster unit then that cluster unit is considered to have failed. The range is 1 to 60 packets.

If the primary unit does not receive a heartbeat packet from a subordinate unit before the heartbeat threshold expires, the primary unit assumes that the subordinate unit has failed.

If a subordinate unit does not receive a heartbeat packet from the primary unit before the heartbeat threshold expires, the subordinate unit assumes that the primary unit has failed. The subordinate unit then begins negotiating to become the new primary unit.

The lower the `hb-lost-threshold` the faster a cluster responds when a unit fails. However, sometimes heartbeat packets may not be sent because a cluster unit is very busy. This can lead to a false positive failure detection. To reduce these false positives you can increase the `hb-lost-threshold`.

Use the following CLI command to increase the lost heartbeat threshold to 12:

```
config system ha
  set hb-lost-threshold 12
end
```

Changing the heartbeat interval

The heartbeat interval is the time between sending HA heartbeat packets. The heartbeat interval range is 1 to 20 (100*ms). The heartbeat interval default is 2 (200 ms).

A heartbeat interval of 2 means the time between heartbeat packets is 200 ms. Changing the heartbeat interval to 5 changes the time between heartbeat packets to 500 ms (5 * 100ms = 500ms).

The HA heartbeat packets consume more bandwidth if the heartbeat interval is short. But if the heartbeat interval is very long, the cluster is not as sensitive to topology and other network changes.

Use the following CLI command to increase the heartbeat interval to 10:

```
config system ha
    set hb-interval 10
end
```

The heartbeat interval combines with the lost heartbeat threshold to set how long a cluster unit waits before assuming that another cluster unit has failed and is no longer sending heartbeat packets. By default, if a cluster unit does not receive a heartbeat packet from a cluster unit for $6 * 200 = 1200$ milliseconds or 1.2 seconds the cluster unit assumes that the other cluster unit has failed.

You can increase both the heartbeat interval and the lost heartbeat threshold to reduce false positives. For example, increasing the heartbeat interval to 20 and the lost heartbeat threshold to 30 means a failure will be assumed if no heartbeat packets are received after $30 * 2000$ milliseconds = 60,000 milliseconds, or 60 seconds.

Use the following CLI command to increase the heartbeat interval to 20 and the lost heartbeat threshold to 30:

```
config system ha
    set hb-lost-threshold 30
    set hb-interval 20
end
```

Changing the time to wait in the hello state

The hello state hold-down time is the number of seconds that a cluster unit waits before changing from hello state to work state. After a failure or when starting up, cluster units operate in the hello state to send and receive heartbeat packets so that all the cluster units can find each other and form a cluster. A cluster unit should change from the hello state to work state after it finds all of the other FortiGate units to form a cluster with. If for some reason all cluster units cannot find each other during the hello state then some cluster units may be joining the cluster after it has formed. This can cause disruptions to the cluster and affect how it operates.

One reason for a delay in all of the cluster units joining the cluster could be the cluster units are located at different sites or if for some other reason communication is delayed between the heartbeat interfaces.

If cluster units are joining your cluster after it has started up or if it takes a while for units to join the cluster you can increase the time that the cluster units wait in the hello state. The hello state hold-down time range is 5 to 300 seconds. The hello state hold-down time default is 20 seconds.

Use the following CLI command to increase the time to wait in the hello state to 1 minute (60 seconds):

```
config system ha
    set hello-holddown 60
end
```

Enabling or disabling HA heartbeat encryption and authentication

You can enable HA heartbeat encryption and authentication to encrypt and authenticate HA heartbeat packets. HA heartbeat packets should be encrypted and authenticated if the cluster interfaces that send HA heartbeat packets are also connected to your networks.

If HA heartbeat packets are not encrypted the cluster password and changes to the cluster configuration could be exposed and an attacker may be able to sniff HA packets to get cluster information. Enabling HA heartbeat message authentication prevents an attacker from creating false HA heartbeat messages. False HA heartbeat messages could affect the stability of the cluster.

HA heartbeat encryption and authentication are disabled by default. Enabling HA encryption and authentication could reduce cluster performance. Use the following CLI command to enable HA heartbeat encryption and authentication.

```
config system ha
```

```
set authentication enable
set encryption enable
end
```

HA authentication and encryption uses AES-128 for encryption and SHA1 for authentication.

Heartbeat bandwidth requirements

The majority of the traffic processed by the HA heartbeat interface is session synchronization traffic. Other heartbeat interface traffic required to synchronize IPsec state/keys, routing tables, configuration changes, and so on is usually negligible.

The amount of traffic required for session synchronization depends on the connections per second (CPS) that the cluster is processing since only new sessions (and session table updates) need to be synchronized.

Another factor to consider is that if session pickup is enabled, traffic on the heartbeat interface surges during a failover or when a unit joins or re-joins the cluster. When one of these events happens, the whole session table needs to be synchronized. Lower bandwidth HA heartbeat interfaces may increase failover time if they can't handle the higher demand during these events.

You can also reduce the amount of heartbeat traffic by:

- Turning off session pickup if you don't need it,
- Configuring `session-pickup-delay` to reduce the number of sessions that are synchronized,
- Using the `session-sync-dev` option to move session synchronization traffic off of the heartbeat link.

See [Session failover \(session-pickup\)](#) on page 262 for details.

Cluster virtual MAC addresses

When a cluster is operating, the FGCP assigns virtual MAC addresses to each primary unit interface. HA uses virtual MAC addresses so that if a failover occurs, the new primary unit interfaces will have the same virtual MAC addresses and IP addresses as the failed primary unit. As a result, most network equipment would identify the new primary unit as the exact same device as the failed primary unit.

If the MAC addresses changed after a failover, the network would take longer to recover because all attached network devices would have to learn the new MAC addresses before they could communicate with the cluster.

If a cluster is operating in NAT/Route mode, the FGCP assigns a different virtual MAC address to each primary unit interface. VLAN subinterfaces are assigned the same virtual MAC address as the physical interface that the VLAN subinterface is added to. Redundant interfaces or 802.3ad aggregate interfaces are assigned the virtual MAC address of the first interface in the redundant or aggregate list.

If a cluster is operating in Transparent mode, the FGCP assigns a virtual MAC address for the primary unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.



A MAC address conflict can occur if two clusters are operating on the same network. See [Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain](#) on page 229 for more information.



Subordinate unit MAC addresses do not change. You can verify this by connecting to the subordinate unit CLI and using the `get hardware interface nic` command to display the MAC addresses of each FortiGate interface.



The MAC address of a reserved management interface is not changed to a virtual MAC address. Instead the reserved management interface keeps its original MAC address.

When the new primary unit is selected after a failover, the primary unit sends gratuitous ARP packets to update the devices connected to the cluster interfaces (usually layer-2 switches) with the virtual MAC address. Gratuitous ARP packets configure connected network devices to associate the cluster virtual MAC addresses and cluster IP address with primary unit physical interfaces and with the layer-2 switch physical interfaces. This is sometimes called using gratuitous ARP packets (sometimes called GARP packets) to train the network. The gratuitous ARP packets sent from the primary unit are intended to make sure that the layer-2 switch forwarding databases (FDBs) are updated as quickly as possible.

Sending gratuitous ARP packets is not required for routers and hosts on the network because the new primary unit will have the same MAC and IP addresses as the failed primary unit. However, since the new primary unit interfaces are connected to different switch interfaces than the failed primary unit, many network switches will update their FDBs more quickly after a failover if the new primary unit sends gratuitous ARP packets.

Changing how the primary unit sends gratuitous ARP packets after a failover

When a failover occurs it is important that the devices connected to the primary unit update their FDBs as quickly as possible to reestablish traffic forwarding.

Depending on your network configuration, you may be able to change the number of gratuitous ARP packets and the time interval between ARP packets to reduce the cluster failover time.

You cannot disable sending gratuitous ARP packets, but you can use the following command to change the number of packets that are sent. For example, enter the following command to send 20 gratuitous ARP packets:

```
config system ha
    set arps 20
end
```

You can use this command to configure the primary unit to send from 1 to 60 ARP packets. Usually you would not change the default setting of 5. In some cases, however, you might want to reduce the number of gratuitous ARP packets. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending a higher number gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully, you could reduce the number of gratuitous ARP packets that are sent to reduce the amount of traffic produced after a failover.

If failover is taking longer than expected, you may be able to reduce the failover time by increasing the number of gratuitous ARP packets sent.

You can also use the following command to change the time interval in seconds between gratuitous ARP packets. For example, enter the following command to change the time between ARP packets to 3 seconds:

```
config system ha
    set arps-interval 3
end
```

The time interval can be in the range of 1 to 20 seconds. The default is 8 seconds between gratuitous ARP packets. Normally you would not need to change the time interval. However, you could decrease the time to be able send more packets in less time if your cluster takes a long time to failover.

There may also be a number of reasons to set the interval higher. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully you could increase the interval to reduce the amount of traffic produced after a failover.

For more information about gratuitous ARP packets see [RFC 826](#) and [RFC 3927](#).

Disabling gratuitous ARP packets after a failover

You can use the following command to turn off sending gratuitous ARP packets after a failover:

```
config system ha
    set gratuitous-arps disable
end
```

Sending gratuitous ARP packets is turned on by default.

In most cases you would want to send gratuitous ARP packets because its a reliable way for the cluster to notify the network to send traffic to the new primary unit. However, in some cases, sending gratuitous ARP packets may be less optimal. For example, if you have a cluster of FortiGates in Transparent mode, after a failover the new primary unit will send gratuitous ARP packets to all of the addresses in its Forwarding Database (FDB). If the FDB has a large number of addresses it may take extra time to send all the packets and the sudden burst of traffic could disrupt the network.

If you choose to disable sending gratuitous ARP packets you must first enable the `link-failed-signal` setting. The cluster must have some way of informing attached network devices that a failover has occurred.

For more information about the `link-failed-signal` setting, see [Updating MAC forwarding tables when a link failover occurs on page 247](#).

How the virtual MAC address is determined

The virtual MAC address is determined based on following formula:

$$00-09-0f-09- <group-id_hex> - (<vcluster_integer> + <idx>)$$

where

`<group-id_hex>` is the HA Group ID for the cluster converted to hexadecimal. The following table lists the virtual MAC address set for each group ID.

HA group ID in integer and hexadecimal format

Integer Group ID	Hexadecimal Group ID
0	00
1	01
2	02

Integer Group ID	Hexadecimal Group ID
3	03
4	04
...	...
10	0a
11	0b
...	...
63	3f
...	...
255	ff

`<vcluster_integer>` is 0 for virtual cluster 1 and 20 for virtual cluster 2. If virtual domains are not enabled, HA sets the virtual cluster to 1 and by default all interfaces are in the root virtual domain. Including virtual cluster and virtual domain factors in the virtual MAC address formula means that the same formula can be used whether or not virtual domains and virtual clustering is enabled.

`<idx>` is the index number of the interface. Interfaces are numbered from 0 to x (where x is the number of interfaces). Interfaces are numbered according to their has map order. See [Interface index and display order on page 220](#). The first interface has an index of 0. The second interface in the list has an index of 1 and so on.



Only the `<idx>` part of the virtual MAC address is different for each interface. The `<vcluster_integer>` would be different for different interfaces if multiple VDOMs have been added.



Between FortiOS releases interface indexing may change so the virtual MAC addresses assigned to individual FortiGate interfaces may also change.

Example virtual MAC addresses

An HA cluster with HA group ID unchanged (default=0) and virtual domains not enabled would have the following virtual MAC addresses for interfaces port1 to port12:

- port1 virtual MAC: 00-09-0f-09-00-00
- port10 virtual MAC: 00-09-0f-09-00-01
- port2 virtual MAC: 00-09-0f-09-00-02
- port3 virtual MAC: 00-09-0f-09-00-03
- port4 virtual MAC: 00-09-0f-09-00-04
- port5 virtual MAC: 00-09-0f-09-00-05

- port6 virtual MAC: 00-09-0f-09-00-06
- port7 virtual MAC: 00-09-0f-09-00-07
- port8 virtual MAC: 00-09-0f-09-00-08
- port9 virtual MAC: 00-09-0f-09-00-
- port11 virtual MAC: 00-09-0f-09-00-0a
- port12 virtual MAC: 00-09-0f-09-00-0b

If the group ID is changed to 34 these virtual MAC addresses change to:

- port1 virtual MAC: 00-09-0f-09-22-00
- port3 virtual MAC: 00-09-0f-09-22-03
- port4 virtual MAC: 00-09-0f-09-22-04
- port5 virtual MAC: 00-09-0f-09-22-05
- port6 virtual MAC: 00-09-0f-09-22-06
- port7 virtual MAC: 00-09-0f-09-22-07
- port8 virtual MAC: 00-09-0f-09-22-08
- port9 virtual MAC: 00-09-0f-09-22-
- port11 virtual MAC: 00-09-0f-09-22-0a
- port12 virtual MAC: 00-09-0f-09-22-0b
- port10 virtual MAC: 00-09-0f-09-22-01
- port2 virtual MAC: 00-09-0f-09-22-02

A cluster with virtual domains enabled where the HA group ID has been changed to 23, port5 and port 6 are in the root virtual domain (which is in virtual cluster1), and port7 and port8 are in the vdom_1 virtual domain (which is in virtual cluster 2) would have the following virtual MAC addresses:

- port5 interface virtual MAC: 00-09-0f-09-23-05
- port6 interface virtual MAC: 00-09-0f-09-23-06
- port7 interface virtual MAC: 00-09-0f-09-23-27
- port8 interface virtual MAC: 00-09-0f-09-23-28

Displaying the virtual MAC address

Every FortiGate physical interface has two MAC addresses: the current hardware address and the permanent hardware address. The permanent hardware address cannot be changed, it is the actual MAC address of the interface hardware. The current hardware address can be changed. The current hardware address is the address seen by the network. For a FortiGate not operating in HA, you can use the following command to change the current hardware address of the port1 interface:

```
config system interface
edit port1
set macaddr <mac_address>
end
end
```

For an operating cluster, the current hardware address of each cluster unit interface is changed to the HA virtual MAC address by the FGCP. The `macaddr` option is not available for a functioning cluster. You cannot change an interface MAC address and you cannot view MAC addresses from the `system interface` CLI command.

You can use the `get hardware nic <interface_name_str>` command to display both MAC addresses for any FortiGate interface. This command displays hardware information for the specified interface. Depending

on their hardware configuration, this command may display different information for different interfaces. You can use this command to display the current hardware address as `Current_HWaddr` and the permanent hardware address as `Permanent_HWaddr`. For some interfaces the current hardware address is displayed as `MAC`. The command displays a great deal of information about the interface so you may have to scroll the output to find the hardware addresses.



You can also use the `diagnose hardware deviceinfo nic <interface_str>` command to display both MAC addresses for any FortiGate interface.

Before HA configuration the current and permanent hardware addresses are the same. For example for one of the units in `Cluster_1`:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 02:09:0f:78:18:c9
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

During HA operation the current hardware address becomes the HA virtual MAC address, for example for the units in `Cluster_1`:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

The following command output for `Cluster_2` shows the same current hardware address for port1 as for the internal interface of `Cluster_2`, indicating a MAC address conflict.

```
FG300A2904500238 # get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 00:09:0F:85:40:FD
.
.
.
```

Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain

A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by

changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID.

This section describes a topology that can result in packet loss, how to determine if packets are being lost, and how to correct the problem by changing the HA Group ID.



Packet loss on a network can also be caused by IP address conflicts. Finding and fixing IP address conflicts can be difficult. However, if you are experiencing packet loss and your network contains two FortiGate HA clusters you can use the information in this article to eliminate one possible source of packet loss.

Changing the HA group ID to avoid MAC address conflicts

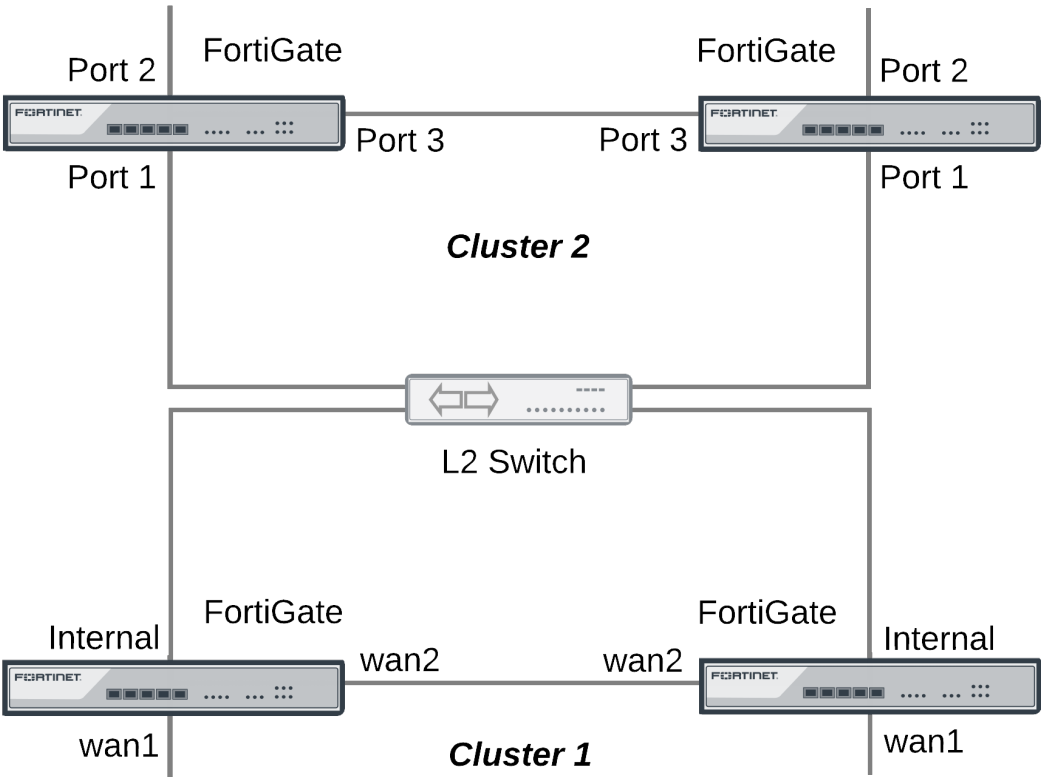
Change the Group ID to change the virtual MAC address of all cluster interfaces. You can change the Group ID from the FortiGate CLI using the following command:

```
config system ha
    set group-id <id_integer>
end
```

Example topology

The topology below shows two clusters. The Cluster_1 internal interfaces and the Cluster_2 port 1 interfaces are both connected to the same broadcast domain. In this topology the broadcast domain could be an internal network. Both clusters could also be connected to the Internet or to different networks.

Example HA topology with possible MAC address conflicts



Ping testing for packet loss

If the network is experiencing packet loss, it is possible that you will not notice a problem unless you are constantly pinging both HA clusters. During normal operation of the network you also might not notice packet loss because the loss rate may not be severe enough to timeout TCP sessions. Also many common types of TCP traffic, such as web browsing, may not be greatly affected by packet loss. However, packet loss can have a significant effect on real time protocols that deliver audio and video data.

To test for packet loss you can set up two constant ping sessions, one to each cluster. If packet loss is occurring the two ping sessions should show alternating replies and timeouts from each cluster.

Cluster_1	Cluster_2
reply	timeout
reply	timeout
reply	timeout
timeout	reply

Cluster_1	Cluster_2
timeout	reply
reply	timeout
reply	timeout
timeout	reply
timeout	reply
timeout	reply
timeout	reply

Viewing MAC address conflicts on attached switches

If two HA clusters with the same virtual MAC address are connected to the same broadcast domain (L2 switch or hub), the MAC address will conflict and bounce between the two clusters. This example Cisco switch MAC address table shows the MAC address flapping between different interfaces (1/0/1 and 1/0/4).

```
1 0009.0f09.0002 DYNAMIC Gi1/0/1
1 0009.0f09.0002 DYNAMIC Gi1/0/4
```

Synchronizing the configuration

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit.

The following settings are not synchronized between cluster units:

- HA override.
- HA device priority.
- The virtual cluster priority.
- The FortiGate host name.
- The HA priority setting for a ping server (or dead gateway detection) configuration.
- The system interface settings of the HA reserved management interface.
- The HA default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.

The primary unit synchronizes all other configuration settings, including the other HA configuration settings.

Disabling automatic configuration synchronization

In some cases you may want to use the following command to disable automatic synchronization of the primary unit configuration to all cluster units.

```
config system ha
    set sync-config disable
end
```


When this option is disabled the cluster no longer synchronizes configuration changes. If a device failure occurs, the new primary unit may not have the same configuration as the failed primary unit. As a result, the new primary unit may process sessions differently or may not function on the network in the same way.

In most cases you should not disable automatic configuration synchronization. However, if you have disabled this feature you can use the `execute ha synchronize` command to manually synchronize a subordinate unit's configuration to that of the primary unit.

You must enter `execute ha synchronize` commands from the subordinate unit that you want to synchronize with the primary unit. Use the `execute ha manage` command to access a subordinate unit CLI.

For example, to access the first subordinate unit and force a synchronization at any time, even if automatic synchronization is disabled enter:

```
execute ha manage 0
execute ha synchronize start
```

You can use the following command to stop a synchronization that is in progress.

```
execute ha synchronize stop
```

Incremental synchronization

When you log into the cluster GUI or CLI to make configuration changes, you are actually logging into the primary unit. All of your configuration changes are first made to the primary unit. Incremental synchronization then immediately synchronizes these changes to all of the subordinate units.

When you log into a subordinate unit CLI (for example using `execute ha manage`) all of the configuration changes that you make to the subordinate unit are also immediately synchronized to all cluster units, including the primary unit, using the same process.

Incremental synchronization also synchronizes other dynamic configuration information such as the DHCP server address lease database, routing table updates, IPsec SAs, MAC address tables, and so on. See [FortiGate HA compatibility with DHCP and PPPoE on page 50](#) for more information about DHCP server address lease synchronization and [Synchronizing kernel routing tables on page 239](#) for information about routing table updates.

Whenever a change is made to a cluster unit configuration, incremental synchronization sends the same configuration change to all other cluster units over the HA heartbeat link. An HA synchronization process running on the each cluster unit receives the configuration change and applies it to the cluster unit. The HA synchronization process makes the configuration change by entering a CLI command that appears to be entered by the administrator who made the configuration change in the first place.

Synchronization takes place silently, and no log messages are recorded about the synchronization activity. However, log messages can be recorded by the cluster units when the synchronization process enters CLI commands. You can see these log messages on the subordinate units if you enable event logging and set the minimum severity level to **Information** and then check the event log messages written by the cluster units when you make a configuration change.

You can also see these log messages on the primary unit if you make configuration changes from a subordinate unit.

Periodic synchronization

Incremental synchronization makes sure that as an administrator makes configuration changes, the configurations of all cluster units remain the same. However, a number of factors could cause one or more cluster units to go out of sync with the primary unit. For example, if you add a new unit to a functioning cluster, the

configuration of this new unit will not match the configuration of the other cluster units. Its not practical to use incremental synchronization to change the configuration of the new unit.

Periodic synchronization is a mechanism that looks for synchronization problems and fixes them. Every minute the cluster compares the configuration file checksum of the primary unit with the configuration file checksums of each of the subordinate units. If all subordinate unit checksums are the same as the primary unit checksum, all cluster units are considered synchronized.

If one or more of the subordinate unit checksums is not the same as the primary unit checksum, the subordinate unit configuration is considered out of sync with the primary unit. The checksum of the out of sync subordinate unit is checked again every 15 seconds. This re-checking occurs in case the configurations are out of sync because an incremental configuration sequence has not completed. If the checksums do not match after 5 checks the subordinate unit that is out of sync retrieves the configuration from the primary unit. The subordinate unit then reloads its configuration and resumes operating as a subordinate unit with the same configuration as the primary unit.

The configuration of the subordinate unit is reset in this way because when a subordinate unit configuration gets out of sync with the primary unit configuration there is no efficient way to determine what the configuration differences are and to correct them. Resetting the subordinate unit configuration becomes the most efficient way to resynchronize the subordinate unit.

Synchronization requires that all cluster units run the same FortiOS firmware build. If some cluster units are running different firmware builds, then unstable cluster operation may occur and the cluster units may not be able to synchronize correctly.



Re-installing the firmware build running on the primary unit forces the primary unit to upgrade all cluster units to the same firmware build.

Console messages when configuration synchronization succeeds

When a cluster first forms, or when a new unit is added to a cluster as a subordinate unit, the following messages appear on the CLI console to indicate that the unit joined the cluster and had its configuring synchronized with the primary unit.

```
slave's configuration is not in sync with master's, sequence:0
slave's configuration is not in sync with master's, sequence:1
  slave's configuration is not in sync with master's, sequence:2
  slave's configuration is not in sync with master's, sequence:3
  slave's configuration is not in sync with master's, sequence:4
  slave starts to sync with master
logout all admin users
slave succeeded to sync with master
```

Console messages when configuration synchronization fails

If you connect to the console of a subordinate unit that is out of synchronization with the primary unit, messages similar to the following are displayed.

```
slave is not in sync with master, sequence:0. (type 0x3)
slave is not in sync with master, sequence:1. (type 0x3)
slave is not in sync with master, sequence:2. (type 0x3)
slave is not in sync with master, sequence:3. (type 0x3)
slave is not in sync with master, sequence:4. (type 0x3)
```

global compared not matched

If synchronization problems occur the console message sequence may be repeated over and over again. The messages all include a type value (in the example `type 0x3`). The type value can help Fortinet Support diagnose the synchronization problem.

HA out of sync object messages and the configuration objects that they reference

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_CONFIGURATION = 0x03	/data/config
HA_SYNC_SETTING_AV = 0x10	
HA_SYNC_SETTING_VIR_DB = 0x11	/etc/vir
HA_SYNC_SETTING_SHARED_LIB = 0x12	/data/lib/libav.so
HA_SYNC_SETTING_SCAN_UNIT = 0x13	/bin/scanunitd
HA_SYNC_SETTING_IMAP_PRXY = 0x14	/bin/imapd
HA_SYNC_SETTING_SMTP_PRXY = 0x15	/bin/smtp
HA_SYNC_SETTING_POP3_PRXY = 0x16	/bin/pop3
HA_SYNC_SETTING_HTTP_PRXY = 0x17	/bin/thttp
HA_SYNC_SETTING_FTP_PRXY = 0x18	/bin/ftpd
HA_SYNC_SETTING_FCNI = 0x19	/etc/fcni.dat
HA_SYNC_SETTING_FDNI = 0x1a	/etc/fdnserver.dat
HA_SYNC_SETTING_FSCI = 0x1b	/etc/sci.dat
HA_SYNC_SETTING_FSAE = 0x1c	/etc/fsae_adgrp.cache
HA_SYNC_SETTING_IDS = 0x20	/etc/ids.rules
HA_SYNC_SETTING_IDSUSER_RULES = 0x21	/etc/idsuser.rules
HA_SYNC_SETTING_IDSCUSTOM = 0x22	
HA_SYNC_SETTING_IDS_MONITOR = 0x23	/bin/ipsmonitor
HA_SYNC_SETTING_IDS_SENSOR = 0x24	/bin/ipsengine
HA_SYNC_SETTING_NIDS_LIB = 0x25	/data/lib/libips.so
HA_SYNC_SETTING_WEBLISTS = 0x30	

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_CONTENTFILTER = 0x31	/data/cmdb/webfilter.bword
HA_SYNC_SETTING_URLFILTER = 0x32	/data/cmdb/webfilter.urlfilter
HA_SYNC_SETTING_FTGD_OVRD = 0x33	/data/cmdb/webfilter.ftgd-ovrd
HA_SYNC_SETTING_FTGD_LRATING = 0x34	/data/cmdb/webfilter.ftgd-ovrd
HA_SYNC_SETTING_EMAILLISTS = 0x40	
HA_SYNC_SETTING_EMAILCONTENT = 0x41	/data/cmdb/spamfilter.bword
HA_SYNC_SETTING_EMAILBWLIST = 0x42	/data/cmdb/spamfilter.emailbwl
HA_SYNC_SETTING_IPBWL = 0x43	/data/cmdb/spamfilter.ipbwl
HA_SYNC_SETTING_MHEADER = 0x44	/data/cmdb/spamfilter.mheader
HA_SYNC_SETTING_RBL = 0x45	/data/cmdb/spamfilter.rbl
HA_SYNC_SETTING_CERT_CONF = 0x50	/etc/cert/cert.conf
HA_SYNC_SETTING_CERT_CA = 0x51	/etc/cert/ca
HA_SYNC_SETTING_CERT_LOCAL = 0x52	/etc/cert/local
HA_SYNC_SETTING_CERT_CRL = 0x53	/etc/cert/crl
HA_SYNC_SETTING_DB_VER = 0x55	
HA_GET_DETAIL_CSUM = 0x71	
HA_SYNC_CC_SIG = 0x75	/etc/cc_sig.dat
HA_SYNC_CC_OP = 0x76	/etc/cc_op
HA_SYNC_CC_MAIN = 0x77	/etc/cc_main
HA_SYNC_FTGD_CAT_LIST = 0x7a	/migadmin/webfilter/ublock/ftgd/data/

Comparing checksums of cluster units

You can use the `diagnose sys ha checksum show` command to compare the configuration checksums of all cluster units. The output of this command shows checksums labelled `global` and `all` as well as checksums for each of the VDOMs including the `root` VDOM. The `get system ha-nonsync-csum` command can be used to display similar information; however, this command is intended to be used by FortiManager.

The primary unit and subordinate unit checksums should be the same. If they are not you can use the `execute ha synchronize start` command to force a synchronization.

The following command output is for the primary unit of a cluster that does not have multiple VDOMs enabled:

```
diagnose sys ha checksum show
is_manage_master()=1, is_root_master()=1
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following command output is for a subordinate unit of the same cluster:

```
diagnose sys ha checksum show
is_manage_master()=0, is_root_master()=0
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following example shows using this command for the primary unit of a cluster with multiple VDOMs. Two VDOMs have been added named `test` and `Eng_vdm`.

From the primary unit:

```
config global
diagnose sys ha checksum show
is_manage_master()=1, is_root_master()=1
debugzone
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53
```

From the subordinate unit:

```
config global
diagnose sys ha checksum show
is_manage_master()=0, is_root_master()=0
debugzone
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
```

```

Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

```

How to diagnose HA out of sync messages

This section describes how to use the `diagnose sys ha checksum show` and `diagnose debug` commands to diagnose the cause of HA out of sync messages.

If HA synchronization is not successful, use the following procedures on each cluster unit to find the cause.

To determine why HA synchronization does not occur

1. Connect to each cluster unit CLI by connected to the console port.
2. Enter the following commands to enable debugging and display HA out of sync messages.

```

diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application hataalk -1
diagnose debug application hasync -1

```

Collect the console output and compare the out of sync messages with the information in the table [HA out of sync object messages and the configuration objects that they reference on page 235](#).

3. Enter the following commands to turn off debugging.

```

diagnose debug disable
diagnose debug reset

```

To determine what part of the configuration is causing the problem

If the previous procedure displays messages that include sync object 0x30 (for example, `HA_SYNC_SETTING_CONFIGURATION = 0x03`) there is a synchronization problem with the configuration. Use the following steps to determine the part of the configuration that is causing the problem.

If your cluster consists of two cluster units, use this procedure to capture the configuration checksums for each unit. If your cluster consists of more that two cluster units, repeat this procedure for all cluster units that returned messages that include 0x30 sync object messages.

1. Connect to each cluster unit CLI by connected to the console port.
2. Enter the following command to turn on terminal capture

```
diagnose debug enable
```

3. Enter the following command to stop HA synchronization.

```
execute ha sync stop
```

4. Enter the following command to display configuration checksums.

```
diagnose sys ha checksum show global
```

5. Copy the output to a text file.
6. Repeat for all affected units.

7. Compare the text file from the primary unit with the text file from each cluster unit to find the checksums that do not match.

You can use a diff function to compare text files.

8. Repeat for the root VDOM:

```
diagnose sys ha checksum show root
```

9. Repeat for all VDOMS (if multiple VDOM configuration is enabled):

```
diagnose sys ha checksum show <vdom-name>
```

10. You can also use the `grep` option to just display checksums for parts of the configuration.

For example to display system related configuration checksums in the root VDOM or log-related checksums in the global configuration:

```
diagnose sys ha checksum root | grep system
diagnose sys ha checksum global | grep log
```

Generally it is the first non-matching checksum that is the cause of the synchronization problem.

11. Attempt to remove/change the part of the configuration that is causing the problem. You can do this by making configuration changes from the primary unit or subordinate unit CLI.

12. Enter the following commands to start HA configuration and stop debugging:

```
execute ha sync start
diagnose debug disable
diagnose debug reset
```

Recalculating the checksums to resolve out of sync messages

Sometimes an error can occur when checksums are being calculated by the cluster. As a result of this calculation error the CLI console could display out of sync error messages even though the cluster is otherwise operating normally. You can also sometimes see checksum calculation errors in `diagnose sys ha checksum` command output when the checksums listed in the `debugzone` output don't match the checksums in the `checksum` part of the output.

One solution to this problem could be to re-calculate the checksums. The re-calculated checksums should match and the out of sync error messages should stop appearing.

You can use the following command to re-calculate HA checksums:

```
diagnose sys ha checksum recalculate [<vdom-name> | global]
```

Just entering the command without options recalculates all checksums. You can specify a VDOM name to just recalculate the checksums for that VDOM. You can also enter `global` to recalculate the global checksum.

Synchronizing kernel routing tables

In a functioning cluster, the primary unit keeps all subordinate unit kernel routing tables (also called the forwarding information base FIB) up to date and synchronized with the primary unit. After a failover, because of these routing table updates the new primary unit does not have to populate its kernel routing table before being able to route traffic. This gives the new primary unit time to rebuild its regular routing table after a failover.

Use the following command to view the regular routing table. This table contains all of the configured routes and routes acquired from dynamic routing protocols and so on. This routing table is not synchronized. On subordinate units this command will not produce the same output as on the primary unit.

```
get router info routing-table
```

Use the following command to view the kernel routing table (FIB). This is the list of resolved routes actually being used by the FortiOS kernel. The output of this command should be the same on the primary unit and the subordinate units.

```
get router info kernel
```

This section describes how clusters handle dynamic routing failover and also describes how to use CLI commands to control the timing of routing table updates of the subordinate unit routing tables from the primary unit.

Controlling how the FGCP synchronizes kernel routing table updates

You can use the following commands to control some of the timing settings that the FGCP uses when synchronizing routing updates from the primary unit to subordinate units and maintaining routes on the primary unit after a failover.

```
config system ha
  set route-hold <hold_integer>
  set route-ttl <ttl_integer>
  set route-wait <wait_integer>
end
```

Change how long routes stay in a cluster unit routing table

Change the `route-ttl` time to control how long routes remain in a cluster unit routing table. The time to live range is 5 to 3600 seconds. The default time to live is 10 seconds.

The time to live controls how long routes remain active in a cluster unit routing table after the cluster unit becomes a primary unit. To maintain communication sessions after a cluster unit becomes a primary unit, routes remain active in the routing table for the route time to live while the new primary unit acquires new routes.

By default, `route-ttl` is set to 10 which may mean that only a few routes will remain in the routing table after a failover. Normally keeping `route-ttl` to 10 or reducing the value to 5 is acceptable because acquiring new routes usually occurs very quickly, especially if graceful restart is enabled, so only a minor delay is caused by acquiring new routes.

If the primary unit needs to acquire a very large number of routes, or if for other reasons, there is a delay in acquiring all routes, the primary unit may not be able to maintain all communication sessions.

You can increase the route time to live if you find that communication sessions are lost after a failover so that the primary unit can use synchronized routes that are already in the routing table, instead of waiting to acquire new routes.

Change the time between routing updates

Change the `route-hold` time to change the time that the primary unit waits between sending routing table updates to subordinate units. The route hold range is 0 to 3600 seconds. The default route hold time is 10 seconds.

To avoid flooding routing table updates to subordinate units, set `route-hold` to a relatively long time to prevent subsequent updates from occurring too quickly. Flooding routing table updates can affect cluster performance if a

great deal of routing information is synchronized between cluster units. Increasing the time between updates means that this data exchange will not have to happen so often.

The `route-hold` time should be coordinated with the `route-wait` time.

Change the time the primary unit waits after receiving a routing update

Change the `route-wait` time to change how long the primary unit waits after receiving routing updates before sending the updates to the subordinate units. For quick routing table updates to occur, set `route-wait` to a relatively short time so that the primary unit does not hold routing table changes for too long before updating the subordinate units.

The `route-wait` range is 0 to 3600 seconds. The default `route-wait` is 0 seconds.

Normally, because the `route-wait` time is 0 seconds the primary unit sends routing table updates to the subordinate units every time its routing table changes.

Once a routing table update is sent, the primary unit waits the `route-hold` time before sending the next update.

Usually routing table updates are periodic and sporadic. Subordinate units should receive these changes as soon as possible so `route-wait` is set to 0 seconds. `route-hold` can be set to a relatively long time because normally the next route update would not occur for a while.

In some cases, routing table updates can occur in bursts. A large burst of routing table updates can occur if a router or a link on a network fails or changes. When a burst of routing table updates occurs, there is a potential that the primary unit could flood the subordinate units with routing table updates. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Setting `route-wait` to a longer time reduces the frequency of additional updates and prevents flooding of routing table updates from occurring.

Configuring graceful restart for dynamic routing failover

When an HA failover occurs, neighbor routers will detect that the cluster has failed and remove it from the network until the routing topology stabilizes. During that time the routers may stop sending IP packets to the cluster and communication sessions that would normally be processed by the cluster may time out or be dropped. Also the new primary unit will not receive routing updates and so will not be able to build and maintain its routing database.

You can solve this problem by configuring graceful restart for the dynamic routing protocols that you are using. This section describes configuring graceful restart for OSPF and BGP.

To support graceful restart you should make sure the new primary unit keeps its synchronized routing data long enough to acquire new routing data. You should also increase the HA route time to live, route wait, and route hold values to 60 using the following CLI command:

```
config system ha
  set route-ttl 60
  set route-wait 60
  set route-hold 60
end
```

Graceful OSPF restart

You can configure graceful restart (also called nonstop forwarding (NSF)) as described in [RFC3623](#) (Graceful OSPF Restart) to solve the problem of dynamic routing failover. If graceful restart is enabled on neighbor routers, they will keep sending packets to the cluster following the HA failover instead of removing it from the network. The neighboring routers assume that the cluster is experiencing a graceful restart.

After the failover, the new primary unit can continue to process communication sessions using the synchronized routing data received from the failed primary unit before the failover. This gives the new primary unit time to update its routing table after the failover.

You can use the following commands to enable graceful restart or NSF on Cisco routers:

```
router ospf 1
 log-adjacency-changes
 nsf ietf helper strict-lsa-checking
```

If the cluster is running OSPF, use the following command to enable graceful restart for OSPF:

```
config router ospf
 set restart-mode graceful-restart
end
```

Graceful BGP restart

If the cluster is running BGP only the primary unit keeps BGP peering connections. When a failover occurs, the BGP peering needs to be reestablished. This will happen if you enable BGP graceful restart which causes the adjacent routers to keep the routes active while the BGP peering is restarted by the new primary unit.



Enabling BGP graceful restart causes the FortiGate's BGP process to restart which can temporarily disrupt traffic through the cluster. So normally you should wait for a quiet time or a maintenance period to enable BGP graceful restart.

Use the following command to enable graceful restart for BGP and set some graceful restart options.

```
config router bgp
 set graceful-restart enable
 set graceful-restart-time 120
 set graceful-stalepath-time 360
 set graceful-update-delay 120
end
```

Notifying BGP neighbors when graceful restart is enabled

You can add BGP neighbors and configure the cluster unit to notify these neighbors that it supports graceful restart.

```
config router bgp
 config neighbor
 edit <neighbor_address_Ipv4>
 set capability-graceful-restart enable
 end
end
```

Bidirectional Forwarding Detection (BFD) enabled BGP graceful restart

You can add a BFD enabled BGP neighbor as a static BFD neighbor using the following command. This example shows how to add a BFD neighbor with IP address 172.20.121.23 that is on the network connected to port4:

```
config router bfd
  config neighbor
    edit 172.20.121.23
      set port4
    end
  end
```

The FGCP supports graceful restart of BFD enabled BGP neighbors. The `config router bfd` command is needed as the BGP auto-start timer is 5 seconds. After HA failover, BGP on the new primary unit has to wait for 5 seconds to connect to its neighbors, and then register BFD requests after establishing the connections. With static BFD neighbors, BFD requests and sessions can be created as soon as possible after the failover. The new command `get router info bfd requests` shows the BFD peer requests.

A BFD session created for a static BFD neighbor/peer request will initialize its state as "INIT" instead of "DOWN" and its detection time `asbfd-required-min-rx * bfd-detect-mult` milliseconds.

When a BFD control packet with nonzero `your_discr` is received, if no session can be found to match the `your_discr`, instead of discarding the packet, other fields in the packet, such as addressing information, are used to choose one session that was just initialized, with zero as its remote discriminator.

When a BFD session in the up state receives a control packet with zero as `your_discr` and down as the state, the session will change its state into down but will not notify this down event to BGP and/or other registered clients.

Link failover (port monitoring or interface monitoring)

Link failover means that if a monitored interface fails, the cluster reorganizes to reestablish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

You configure monitored interfaces (also called interface monitoring or port monitoring) by selecting the interfaces to monitor as part of the cluster HA configuration.

You can monitor up to 64 interfaces.

The interfaces that you can monitor appear on the port monitor list. You can monitor all FortiGate interfaces including redundant interfaces and 802.3ad aggregate interfaces.

You cannot monitor the following types of interfaces (you cannot select the interfaces on the port monitor list):

- FortiGate interfaces that contain an internal switch.
- VLAN subinterfaces.
- IPsec VPN interfaces.
- Individual physical interfaces that have been added to a redundant or 802.3ad aggregate interface.
- FortiGate-5000 series backplane interfaces that have not been configured as network interfaces.

If you are configuring a virtual cluster you can create a different port monitor configuration for each virtual cluster. Usually for each virtual cluster you would monitor the interfaces that have been added to the virtual domains in each virtual cluster.



Wait until after the cluster is up and running to enable interface monitoring. You do not need to configure interface monitoring to get a cluster up and running and interface monitoring will cause failovers if for some reason during initial setup a monitored interface has become disconnected. You can always enable interface monitoring once you have verified that the cluster is connected and operating properly.



You should only monitor interfaces that are connected to networks, because a failover may occur if you monitor an unconnected interface.

To enable interface monitoring - GUI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster GUI.
2. Go to **System > HA** and edit the primary unit (**Role** is **MASTER**).
3. Select the **Port Monitor** check boxes for the **port1** and **port2** interfaces and select **OK**.

The configuration change is synchronized to all cluster units.

To enable interface monitoring - CLI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster CLI.
2. Enter the following command to enable interface monitoring for port1 and port2.

```
configure system ha
  set monitor port1 port2
end
```

The following example shows how to enable monitoring for the external, internal, and DMZ interfaces.

```
config system ha
  set monitor external internal dmz
end
```

With interface monitoring enabled, during cluster operation, the cluster monitors each cluster unit to determine if the monitored interfaces are operating and connected. Each cluster unit can detect a failure of its network interface hardware. Cluster units can also detect if its network interfaces are disconnected from the switch they should be connected to.



Cluster units cannot determine if the switch that its interfaces are connected to is still connected to the network. However, you can use remote IP monitoring to make sure that the cluster unit can connect to downstream network devices. See [Remote link failover on page 250](#).

Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link between a network and the primary unit fails, to maintain communication with this network, the cluster must select a different primary unit; one that is still connected to the network. Unless another link failure has occurred, the new primary unit will have an active link to the network and will be able to maintain communication with it.

To support link failover, each cluster unit stores link state information for all monitored cluster units in a link state database. All cluster units keep this link state database up to date by sharing link state information with the other cluster units. If one of the monitored interfaces on one of the cluster units becomes disconnected or fails, this information is immediately shared with all cluster units.

If a monitored interface on the primary unit fails

If a monitored interface on the primary unit fails, the cluster renegotiates to select a new primary unit using the process described in [Primary unit selection on page 39](#). Because the cluster unit with the failed monitored interface has the lowest monitor priority, a different cluster unit becomes the primary unit. The new primary unit should have fewer link failures.

After the failover, the cluster resumes and maintains communication sessions in the same way as for a device failure. See [Device failover on page 216](#).

If a monitored interface on a subordinate unit fails

If a monitored interface on a subordinate unit fails, this information is shared with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster.

In an active-passive cluster after a subordinate unit link failover, the subordinate unit continues to function normally as a subordinate unit in the cluster.


In an active-active cluster after a subordinate unit link failure:

- The subordinate unit with the failed monitored interface can continue processing connections between functioning interfaces. However, the primary unit stops sending sessions to a subordinate unit that use any failed monitored interfaces on the subordinate unit.
- If session pickup is enabled, all sessions being processed by the subordinate unit failed interface that can be failed over are failed over to other cluster units. Sessions that cannot be failed over are lost and have to be restarted.
- If session pickup is not enabled all sessions being processed by the subordinate unit failed interface are lost.

How link failover maintains traffic flow

Monitoring an interface means that the interface is connected to a high priority network. As a high priority network, the cluster should maintain traffic flow to and from the network, even if a link failure occurs. Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link that the primary unit has to a high priority network fails, to maintain traffic flow to and from this network, the cluster must select a different primary unit. This new primary unit should have an active link to the high priority network.

A link failure causes a cluster to select a new primary unit



If a monitored interface on the primary unit fails, the cluster renegotiates and selects the cluster unit with the highest monitor priority to become the new primary unit. The cluster unit with the highest monitor priority is the cluster unit with the most monitored interfaces connected to networks.

After a link failover, the primary unit processes all traffic and all subordinate units, even the cluster unit with the link failure, share session and link status. In addition all configuration changes, routes, and IPsec SAs are synchronized to the cluster unit with the link failure.

In an active-active cluster, the primary unit load balances traffic to all the units in the cluster. The cluster unit with the link failure can process connections between its functioning interfaces (for, example if the cluster has connections to an internal, external, and DMZ network, the cluster unit with the link failure can still process connections between the external and DMZ networks).

Recovery after a link failover and controlling primary unit selection (controlling falling back to the prior primary unit)

If you find and correct the problem that caused a link failure (for example, re-connect a disconnected network cable) the cluster updates its link state database and re-negotiates to select a primary unit.

What happens next depends on how the cluster configuration affects primary unit selection:

- The former primary unit will once again become the primary unit (falling back to becoming the primary unit)
- The primary unit will not change.

As described in [Primary unit selection and age on page 41](#), when the link is restored, if no options are configured to control primary unit selection and the cluster age difference is less than 300 seconds the former primary unit will once again become the primary unit. If the age differences are greater than 300 seconds then a new primary unit is not selected. Since you have no control on the age difference the outcome can be unpredictable. This is not a problem in cases where its not important which unit becomes the primary unit.

Preventing a primary unit change after a failed link is restored

Some organizations will not want the cluster to change primary units when the link is restored. Instead they would rather wait to restore the primary unit during a maintenance window. This functionality is not directly supported, but you can experiment with changing some primary unit selection settings. For example, in most cases it should work to enable override on all cluster units and make sure their priorities are the same. This should mean that the primary unit should not change after a failed link is restored.

Then, when you want to restore the original primary unit during a maintenance window you can just set its Device Priority higher. After it becomes the primary unit you can reset all device priorities to the same value. Alternatively during a maintenance window you could reboot the current primary unit and any subordinate units except the one that you want to become the primary unit.

If the `override` CLI keyword is enabled on one or more cluster units and the device priority of a cluster unit is set higher than the others, when the link failure is repaired and the cluster unit with the highest device priority will always become the primary unit.

Testing link failover

You can test link failure by disconnecting the network cable from a monitored interface of a cluster unit. If you disconnect a cable from a primary unit monitored interface the cluster should renegotiate and select one of the other cluster units as the primary unit. You can also verify that traffic received by the disconnected interface continues to be processed by the cluster after the failover.

If you disconnect a cable from a subordinate unit interface the cluster will not renegotiate.

Updating MAC forwarding tables when a link failover occurs

When a FortiGate HA cluster is operating and a monitored interface fails on the primary unit, the primary unit usually becomes a subordinate unit and another cluster unit becomes the primary unit. After a link failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables (also called arp tables) of the switches connected to the cluster. This is normal link failover operation.

Even when gratuitous ARP packets are sent, some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur if the switch does not detect the failure and does not clear its MAC forwarding table.

You have another option available to make sure the switch detects the failover and clears its MAC forwarding tables. You can use the following command to cause a cluster unit with a monitored interface link failure to briefly shut down all of its interfaces (except the heartbeat interfaces) after the failover occurs:

```
config system ha
    set link-failed-signal enable
end
```

Usually this means each interface of the former primary unit is shut down for about a second. When this happens the switch should be able to detect this failure and clear its MAC forwarding tables of the MAC addresses of the

former primary unit and pickup the MAC addresses of the new primary unit. Each interface will shut down for a second but the entire process usually takes a few seconds. The more interfaces the FortiGate has, the longer it will take.

Normally, the new primary unit also sends gratuitous ARP packets that also help the switch update its MAC forwarding tables to connect to the new primary unit. If `link-failed-signal` is enabled, sending gratuitous ARP packets is optional and can be disabled if you don't need it or if its causing problems. See [Disabling gratuitous ARP packets after a failover on page 226](#)

Multiple link failures

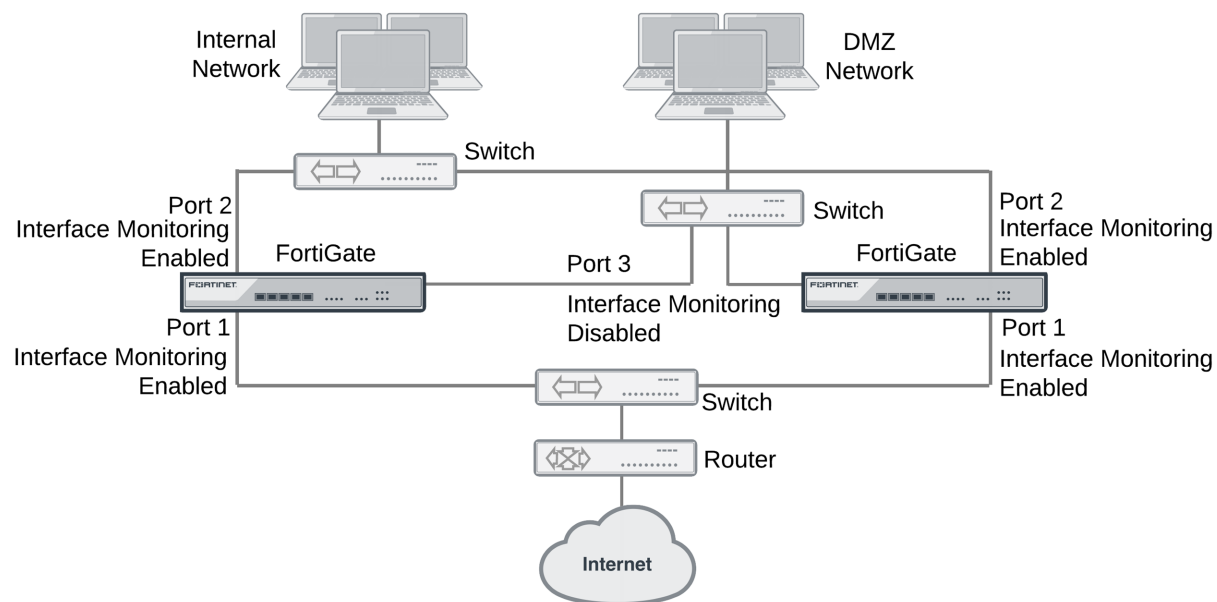
Every time a monitored interface fails, the cluster repeats the processes described above. If multiple monitored interfaces fail on more than one cluster unit, the cluster continues to negotiate to select a primary unit that can provide the most network connections.

Example link failover scenarios

For the following examples, assume a cluster configuration consisting of two FortiGates (FGT_1 and FGT_2) connected to three networks: internal using port2, external using port1, and DMZ using port3. In the HA configuration, the device priority of FGT_1 is set higher than the unit priority of FGT_2.

The cluster processes traffic flowing between the internal and external networks, between the internal and DMZ networks, and between the external and DMZ networks. If there are no link failures, FGT1 becomes the primary unit because it has the highest device priority.

Sample link failover scenario topology



Example the port1 link on FGT_1 fails

If the port1 link on FGT_1 fails, FGT_2 becomes primary unit because it has fewer interfaces with a link failure. If the cluster is operating in active-active mode, the cluster load balances traffic between the internal network (port2) and the DMZ network (port3). Traffic between the Internet (port1) and the internal network (port2) and between the Internet (port1) and the DMZ network (port3) is processed by the primary unit only.

Example port2 on FGT_1 and port1 on FGT_2 fail

If port2 on FGT_1 and port1 on FGT_2 fail, then FGT_1 becomes the primary unit. After both of these link failures, both cluster units have the same monitor priority. So the cluster unit with the highest device priority (FGT_1) becomes the primary unit.

Only traffic between the Internet (port1) and DMZ (port3) networks can pass through the cluster and the traffic is handled by the primary unit only. No load balancing will occur if the cluster is operating in active-active mode.

Monitoring VLAN interfaces

If the FortiGates in the cluster have VLAN interfaces, you can use the following command to monitor all VLAN interfaces and write a log message if one of the VLAN interfaces is found to be down.

Once configured, this feature works by verifying that the primary unit can connect to the subordinate unit over each VLAN. This verifies that the switch that the VLAN interfaces are connected to is configured correctly for each VLAN. If the primary unit cannot connect to the subordinate unit over one of the configured VLANs the primary unit writes a link monitor log message indicating that the named VLAN went down (log message id 20099).

Use the following CLI command to enable monitoring VLAN interfaces:

```
config system ha-monitor
  set monitor-vlan enable/disable
  set vlan-hb-interval <interval_seconds>
  set vlan-hb-lost-threshold <vlan-lost-heartbeat-threshold>
end
```

`vlan-hb-interval` is the time between sending VLAN heartbeat packets over the VLAN. The VLAN heartbeat range is 1 to 30 seconds. The default is 5 seconds.

`vlan-hb-lost-threshold` is the number of consecutive VLAN heartbeat packets that are not successfully received across the VLAN before assuming that the VLAN is down. The default value is 3, meaning that if 3 heartbeat packets sent over the VLAN are not received then the VLAN is considered to be down. The range is 1 to 60 packets.

A VLAN heartbeat interval of 5 means the time between heartbeat packets is five seconds. A VLAN heartbeat threshold of 3 means it takes $5 \times 3 = 15$ seconds to detect that a VLAN is down.

Subsecond failover

On FortiGate models 395xB and 3x40B HA link failover supports subsecond failover (that is a failover time of less than one second). Subsecond failover is available for interfaces that can issue a link failure system call when the interface goes down. When an interface experiences a link failure and sends the link failure system call, the FGCP receives the system call and initiates a link failover.

For interfaces that do not support subsecond failover, port monitoring regularly polls the connection status of monitored interfaces. When a check finds that an interface has gone down, port monitoring causes a link failover. Subsecond failover results in a link failure being detected sooner because the system doesn't have to wait for the next poll to find out about the failure.

Subsecond failover can accelerate HA failover to reduce the link failover time to less than one second under ideal conditions. Actual failover performance may vary depending on traffic patterns and network configuration. For example, some network devices may respond slowly to an HA failover.

No configuration changes are required to support subsecond failover. However, for best subsecond failover results, the recommended heartbeat interval is 100ms and the recommended lost heartbeat threshold is 5 (see [Modifying heartbeat timing on page 222](#)).

```
config system ha
    set hb-lost-threshold 5
    set hb-interval 1
end
```

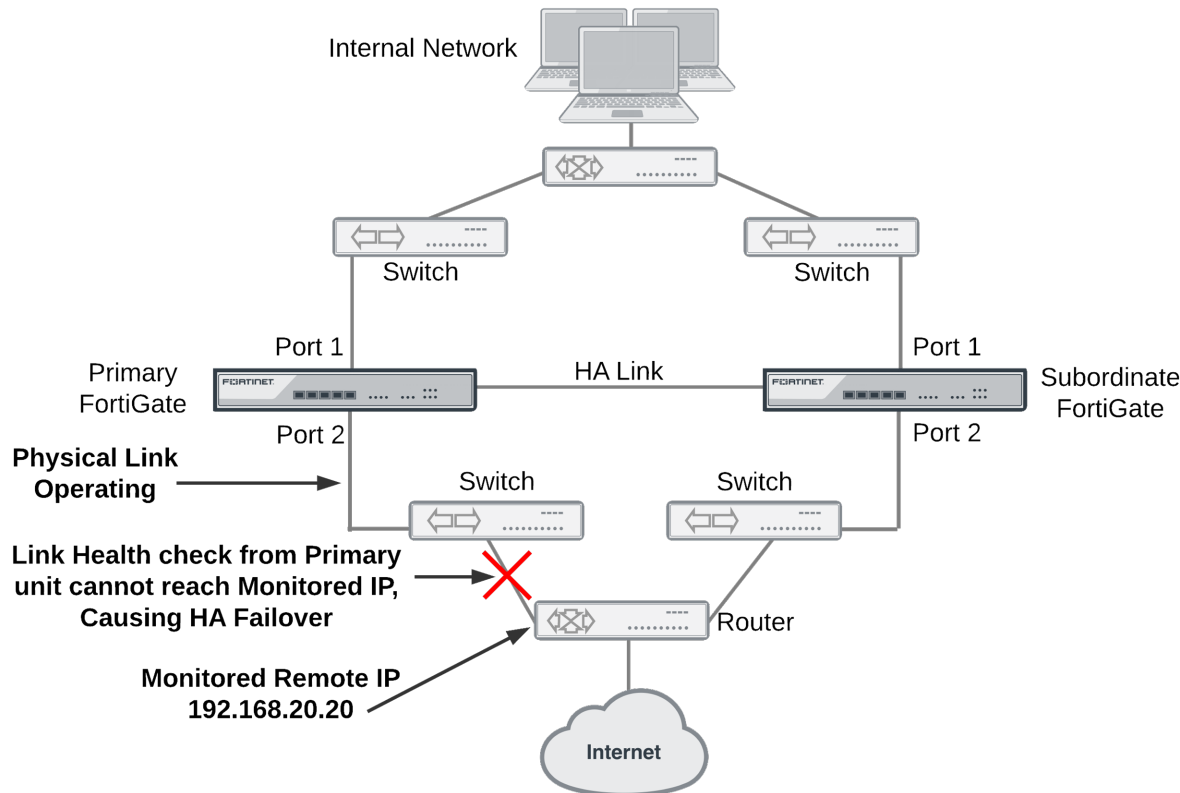
For information about how to reduce failover times, see [Failover performance on page 260](#).

Remote link failover

Remote link failover (also called remote IP monitoring) is similar to HA port monitoring and link health monitoring (also known as dead gateway detection). Port monitoring causes a cluster to failover if a monitored primary unit interface fails or is disconnected. Remote IP monitoring uses link health monitors configured for FortiGate interfaces on the primary unit to test connectivity with IP addresses of network devices. Usually these would be IP addresses of network devices not directly connected to the cluster. For example, a downstream router. Remote IP monitoring causes a failover if one or more of these remote IP addresses does not respond to link health checking.

By being able to detect failures in network equipment not directly connected to the cluster, remote IP monitoring can be useful in a number of ways depending on your network configuration. For example, in a full mesh HA configuration, with remote IP monitoring, the cluster can detect failures in network equipment that is not directly connected to the cluster but that would interrupt traffic processed by the cluster if the equipment failed.

Example HA remote IP monitoring topology



In the simplified example topology shown above, the switch connected directly to the primary unit is operating normally but the link on the other side of the switches fails. As a result traffic can no longer flow between the primary unit and the Internet.

To detect this failure you can create a link health monitor for port2 that causes the primary unit to test connectivity to 192.168.20.20. If the health monitor cannot connect to 192.268.20.20 the cluster fails over and the subordinate unit becomes the new primary unit. After the failover, the health check monitor on the new primary unit can connect to 192.168.20.20 so the failover maintains connectivity between the internal network and the Internet through the cluster.

To configure remote IP monitoring

1. Enter the following commands to configure HA remote monitoring for the example topology.
 - Enter the `pingserver-monitor-interface` keyword to enable HA remote IP monitoring on port2.
 - Leave the `pingserver-failover-threshold` set to the default value of 5. This means a failover occurs if the link health monitor doesn't get a response after 5 attempts.
 - Enter the `pingserver-flip-timeout` keyword to set the flip timeout to 120 minutes. After a failover, if HA remote IP monitoring on the new primary unit also causes a failover, the flip timeout prevents the failover from occurring until the timer runs out. Setting the `pingserver-flip-timeout` to 120 means that remote IP monitoring can only cause a failover every 120 minutes. This flip timeout is required to prevent repeating failovers if remote IP monitoring causes a failover from all cluster units because none of the cluster units can connect to the monitored IP addresses.

```
config system ha
```

```

set pingserver-monitor-interface port2
set pingserver-failover-threshold 5
set pingserver-flip-timeout 120
end

```

2. Enter the following commands to add a link health monitor for the port2 interface and to set HA remote IP monitoring priority for this link health monitor.

- Enter the `detectserver` keyword to set the health monitor server IP address to 192.168.20.20.
- Leave the `ha-priority` keyword set to the default value of 1. You only need to change this priority if you change the HA `pingserver-failover-threshold`. The `ha-priority` setting is not synchronized among cluster units.



The `ha-priority` setting is not synchronized among cluster units. So if you want to change the `ha-priority` setting you must change it separately on each cluster unit. Otherwise it will remain set to the default value of 1.

- Use the `interval` keyword to set the time between link health checks and use the `failtime` keyword to set the number of times that a health check can fail before a failure is detected (the failover threshold). The following example reduces the failover threshold to 2 but keeps the health check interval at the default value of 5.

```

config system link-monitor
edit ha-link-monitor
set server 192.168.20.20
set srcintf port2
set ha-priority 1
set interval 5
set failtime 2
end

```

Adding HA remote IP monitoring to multiple interfaces

You can enable HA remote IP monitoring on multiple interfaces by adding more interface names to the `pingserver-monitor-interface` keyword. If your FortiGate configuration includes VLAN interfaces, aggregate interfaces and other interface types, you can add the names of these interfaces to the `pingserver-monitor-interface` keyword to configure HA remote IP monitoring for these interfaces.

For example, enable remote IP monitoring for interfaces named port2, port20, and vlan_234:

```

config system ha
set pingserver-monitor-interface port2 port20 vlan_234
set pingserver-failover-threshold 10
set pingserver-flip-timeout 120
end

```

Then configure health monitors for each of these interfaces. In the following example, default values are accepted for all settings other than the server IP address.

```

config system link-monitor
edit port2
set server 192.168.20.20
next
edit port20
set server 192.168.20.30
next
edit vlan_234

```

```

    set server 172.20.12.10
end

```

Changing the link monitor failover threshold

If you have multiple link monitors you may want a failover to occur only if more than one of them fails.

For example, you may have 3 link monitors configured on three interfaces but only want a failover to occur if two of the link monitors fail. To do this you must set the HA priorities of the link monitors and the HA `pingserver-failover-threshold` so that the priority of one link monitor is less than the failover threshold but the added priorities of two link monitors is equal to or greater than the failover threshold. Failover occurs when the HA priority of all failed link monitors reaches or exceeds the threshold.

For example, set the failover threshold to 10 and monitor three interfaces:

```

config system ha
    set pingserver-monitor-interface port2 port20 vlan_234
    set pingserver-failover-threshold 10
    set pingserver-flip-timeout 120
end

```

Then set the HA priority of link monitor server to 5.



The HA Priority (`ha-priority`) setting is not synchronized among cluster units. In the following example, you must set the HA priority to 5 by logging into each cluster unit.

```

config system link-monitor
    edit port2
        set srcintf port2
        set server 192.168.20.20
        set ha-priority 5
    next
    edit port20
        set srcintf port20
        set server 192.168.20.30
        set ha-priority 5
    next
    edit vlan_234
        set srcintf vlan_234
        set server 172.20.12.10
        set ha-priority 5
    end
end

```

If only one of the link monitors fails, the total link monitor HA priority will be 5, which is lower than the failover threshold so a failover will not occur. If a second link monitor fails, the total link monitor HA priority of 10 will equal the failover threshold, causing a failover.

By adding multiple link monitors and setting the HA priorities for each, you can fine tune remote IP monitoring. For example, if it is more important to maintain connections to some networks you can set the HA priorities higher for these link monitors. And if it is less important to maintain connections to other networks you can set the HA priorities lower for these link monitors. You can also adjust the failover threshold so that if the cluster cannot connect to one or two high priority IP addresses a failover occurs. But a failover will not occur if the cluster cannot connect to one or two low priority IP addresses.

Monitoring multiple IP addresses from one interface

You can add multiple IP addresses to a single link monitor to use HA remote IP monitoring to monitor more than one IP address from a single interface. If you add multiple IP addresses, the health checking will be with all of the addresses at the same time. The link monitor only fails when no responses are received from all of the addresses.

```
config system link-monitor
  edit port2
    set srcintf port2
    set server 192.168.20.20 192.168.20.30 172.20.12.10
  end
```

Flip timeout

The HA remote IP monitoring configuration also involves setting a flip timeout. The flip timeout is required to reduce the frequency of failovers if, after a failover, HA remote IP monitoring on the new primary unit also causes a failover. This can happen if the new primary unit cannot connect to one or more of the monitored remote IP addresses. The result could be that until you fix the network problem that blocks connections to the remote IP addresses, the cluster will experience repeated failovers. You can control how often the failovers occur by setting the flip timeout. The flip timeout stops HA remote IP monitoring from causing a failover until the primary unit has been operating for the duration of the flip timeout.

If you set the flip timeout to a relatively high number of minutes you can find and repair the network problem that prevented the cluster from connecting to the remote IP address without the cluster experiencing very many failovers. Even if it takes a while to detect the problem, repeated failovers at relatively long time intervals do not usually disrupt network traffic.

Use the following command to set the flip timeout to 3 hours (360 minutes):

```
config system ha
  set pingserver-flip-timeout 360
end
```

Restoring normal cluster operation after the remote link is restored

In a remote IP monitoring configuration, if you also want the same cluster unit to always be the primary unit you can set its device priority higher and enable override. With this configuration, when a remote IP monitoring failover occurs, after the flip timeout expires another failover will occur (because override is enabled) and the unit with override enabled becomes the primary unit again. So the cluster automatically returns to normal operation.

The primary unit starts remote IP monitoring again. If the remote link is restored the cluster continues to operate normally. If, however, the remote link is still down, remote link failover causes the cluster to failover again. This will repeat each time the flip timeout expires until the failed remote link is restored.

You can use the `pingserver-slave-force-reset` option to control this behavior. By default this option is enabled and the behavior described above occurs. The overall behavior is that when the remote link is restored the cluster automatically returns to normal operation after the flip timeout.

If you disable `pingserver-slave-force-reset` after the initial remote IP monitoring failover nothing will happen after the flip timeout (as long as the new primary unit doesn't experience some kind of failover). The result is that repeated failovers no longer happen. But it also means that the original primary unit will remain the subordinate unit and will not resume operating as the primary unit.

Detecting HA remote IP monitoring failovers

Just as with any HA failover, you can detect HA remote IP monitoring failovers by using SNMP to monitor for HA traps. You can also use alert email to receive notifications of HA status changes and monitor log messages for HA failover log messages. In addition, FortiGates send the critical log message `Ping Server is down` when a ping server fails. The log message includes the name of the interface that the ping server has been added to.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time experienced by your network users may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Monitoring cluster units for failover

You can use logging and SNMP to monitor cluster units for failover. Both the primary and subordinate units can be configured to write log messages and send SNMP traps if a failover occurs. You can also log into the cluster GUI and CLI to determine if a failover has occurred.

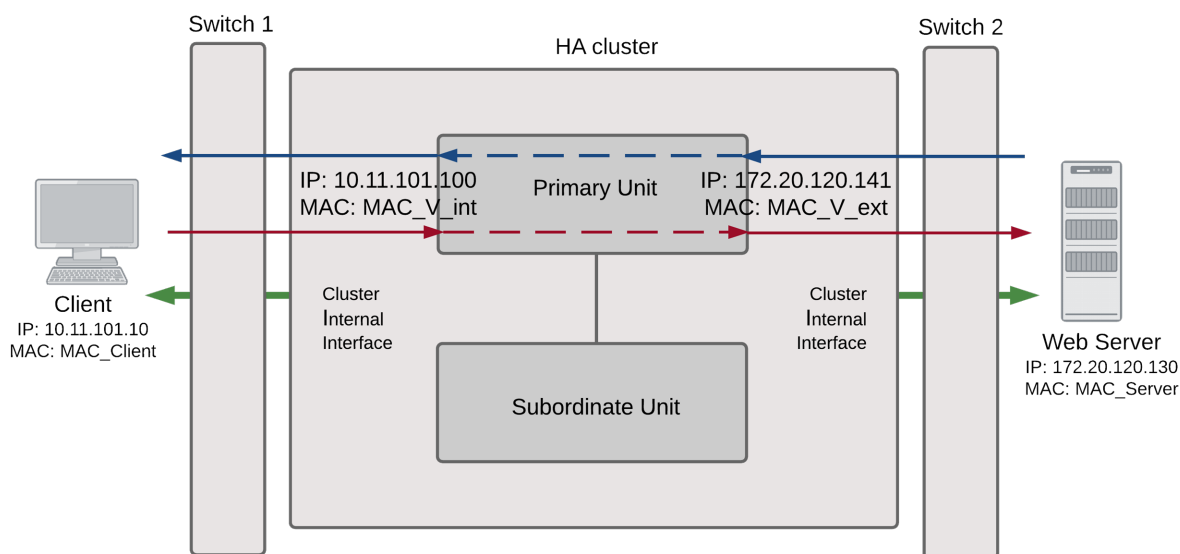
NAT/Route mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In an active-passive cluster operating in NAT/Route mode, four MAC addresses are involved in communication between the client and the web server when the primary unit processes the connection:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only know the cluster external virtual MAC address (MAC_V_ext).

NAT/Route mode active-passive cluster packet flow**Packet flow from client to web server**

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

6. The primary unit processes the packet.
7. The primary unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_V_ext
Destination	172.20.120.130	MAC_Server

8. The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from web server to client

1. When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
2. The web server issues an ARP request to 172.20.120.141.
3. The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.
4. The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

5. The primary unit processes the packet.
6. The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_V_int
Destination	10.11.101.10	MAC_Client

7. The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails the subordinate unit becomes the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
3. The new primary unit sends gratuitous ARP packets from the internal interface to the 10.11.101.0 network to associate its internal IP address with the internal virtual MAC address.
4. The new primary unit sends gratuitous ARP packets to the 172.20.120.0 to associate its external IP address with the external virtual MAC address.
5. Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Transparent mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.

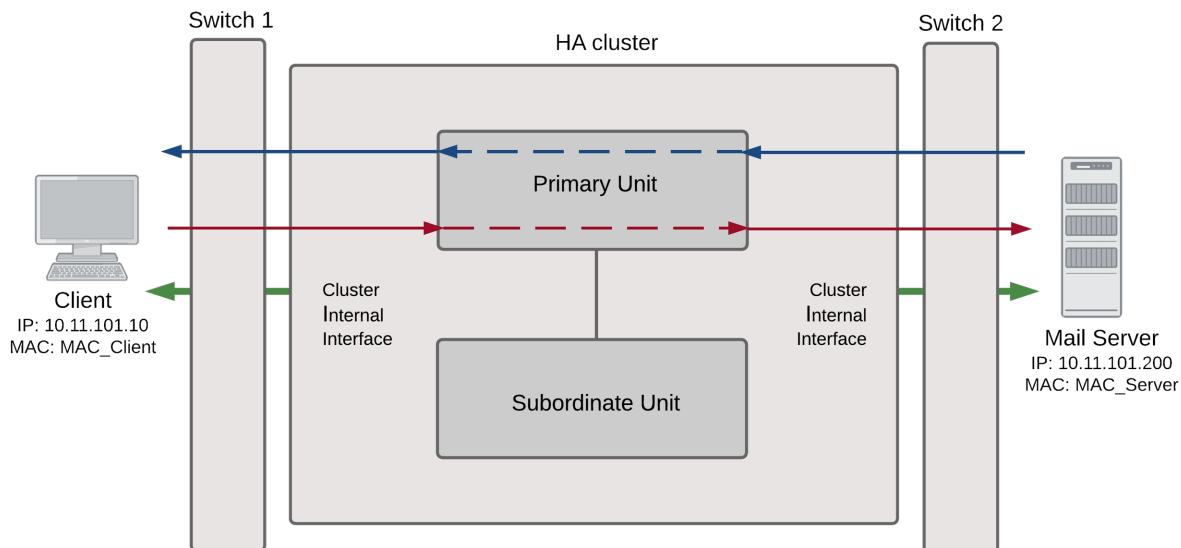
In an active-passive cluster operating in Transparent mode, two MAC addresses are involved in the communication between a client and a server when the primary unit processes a connection:

- Client MAC address (MAC_Client)
- Server MAC address (MAC_Server)

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and processed by the cluster.

The cluster's presence on the network is transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the HA virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the HA virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Transparent mode active-passive packet flow



Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- The primary unit processes the packet.
- The primary unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from mail server to client

- To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
- The primary unit forwards the ARP request to the client computer.
- The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
- The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- The primary unit processes the packet.
- The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- If the primary unit fails, the subordinate unit negotiates to become the primary unit.
- The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
- The new primary unit sends gratuitous ARP packets to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.

4. The new primary unit sends gratuitous ARP packets to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
5. Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Failover performance

This section describes the designed device and link failover times for a FortiGate cluster and also shows results of a failover performance test.

Device failover performance

By design FGCP device failover time is 2 seconds for a two-member cluster with ideal network and traffic conditions. If subsecond failover is enabled the failover time can drop below 1 second.

All cluster units regularly receive HA heartbeat packets from all other cluster units over the HA heartbeat link. If any cluster unit does not receive a heartbeat packet from any other cluster unit for 2 seconds, the cluster unit that has not sent heartbeat packets is considered to have failed.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions. Typically if subsecond failover is not enabled you can expect a failover time of 9 to 15 seconds depending on the cluster and network configuration. The failover time can also be increased by more complex configurations and or configurations with network equipment that is slow to respond.

You can change the `hb-lost-threshold` to increase or decrease the device failover time. See [Modifying heartbeat timing on page 222](#) for information about using `hb-lost-threshold`, and other heartbeat timing settings.

Link failover performance

Link failover time is controlled by how long it takes for a cluster to synchronize the cluster link database. When a link failure occurs, the cluster unit that experienced the link failure uses HA heartbeat packets to broadcast the updated link database to all cluster units. When all cluster units have received the updated database the failover is complete.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions.

Reducing failover times

- Keep the network configuration as simple as possible with as few as possible network connections to the cluster.
- If possible operate the cluster in Transparent mode.
- Use high-performance switches so that the switches failover to interfaces connected to the new primary unit as quickly as possible.
- Use accelerated FortiGate interfaces. In some cases accelerated interfaces will reduce failover times.
- Make sure the FortiGate sends multiple gratuitous arp packets after a failover. In some cases, sending more gratuitous arp packets will cause connected network equipment to recognize the failover sooner. To send 10 gratuitous arp packets:

```
config system ha
```

```
    set arps 10
end
```

- Reduce the time between gratuitous arp packets. This may also caused connected network equipment to recognize the failover sooner. To send 50 gratuitous arp packets with 1 second between each packet:

```
config system ha
    set arps 50
    set arps-interval 1
end
```

- Reduce the number of lost heartbeat packets and reduce the heartbeat interval timers to be able to more quickly detect a device failure. To set the lost heartbeat threshold to 3 packets and the heartbeat interval to 100 milliseconds:

```
config system ha
    set hb-interval 1
    set hb-lost-threshold 3
end
```

- Reduce the hello state hold down time to reduce the amount of the time the cluster waits before transitioning from the hello to the work state. To set the hello state hold down time to 5 seconds:

```
config system ha
    set hello-holddown 5
end
```

- Enable sending a link failed signal after a link failover to make sure that attached network equipment responds as quickly as possible to a link failure. To enable the link failed signal:

```
config system ha
    set link-failed-signal enable
end
```

Session failover (session-pickup)

Session failover means that after the primary unit fails, communications sessions resume on the new primary unit with minimal or no interruption. Two categories of sessions need to be resumed after a failover:

- Sessions passing through the cluster
- Sessions terminated by the cluster

If you enable session failover (also called session-pickup) for the cluster, during cluster operation the primary unit informs the subordinate units of changes to the primary unit connection and state tables for sessions passing through the cluster, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster.

After a failover the new primary unit recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary unit and are handled according to their last known state.



Session-pickup has some limitations. For example, session failover is not supported for sessions being scanned by proxy-based security profiles. Session failover is supported for sessions being scanned by flow-based security profiles; however, flow-based sessions that fail over are not inspected after they fail over. For more limitations, see [Session failover limitations for sessions passing through the cluster on page 264](#).

Sessions terminated by the cluster include management sessions (such as HTTPS connections to the FortiGate GUI or SSH connection to the CLI as well as SNMP and logging and so on). Also included in this category are IPsec VPN, SSL VPN, sessions terminated by the cluster, explicit proxy, WAN Optimization and web caching. In general, whether or not session-pickup is enabled, these sessions do not failover and have to be restarted. There are some exceptions though, particularly for IPsec and SSL VPN. For more information, see [Session failover limitations for sessions terminated by the cluster on page 267](#).

Enabling session-pickup for TCP, UDP and ICMP session failover

To enable session-pickup, go to **System > HA** and enable session-pickup.

From the CLI enter:

```
config system ha
    set session-pickup enable
end
```

When session-pickup is enabled, the FGCP synchronizes the primary unit's TCP session table to all cluster units. As soon as a new TCP session is added to the primary unit session table, that session is synchronized to all cluster units. This synchronization happens as quickly as possible to keep the session tables synchronized.

If the primary unit fails, the new primary unit uses its synchronized session table to resume all TCP sessions that were being processed by the former primary unit with only minimal interruption. Under ideal conditions all TCP sessions should be resumed. This is not guaranteed though and under less than ideal conditions some TCP sessions may need to be restarted.

If session pickup is enabled, you can use the following command to also enable UDP and ICMP session failover:

```
config system ha
  set session-pickup-connectionless enable
end
```



Session pickup does not support multicast session failover.

If session pickup is disabled

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed. Most sessions can be resumed as a normal result of how TCP/IP communications resume communication after any routine network interruption.



The session-pickup setting does not affect session failover for sessions terminated by the cluster.

If you do not require session failover protection, leaving session pickup disabled may reduce CPU usage and reduce HA heartbeat network bandwidth usage. Also if your cluster is mainly being used for traffic that is not synchronized (for example, for proxy-based security profile processing) enabling session pickup is not recommended since most sessions will not be failed over anyway.

If session pickup is not enabled, the FGCP does not synchronize the primary unit session table to other cluster units and sessions do not resume after a failover. After a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates.

Many protocols can successfully restart sessions with little, if any, loss of data. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Since most HTTP sessions are very short, in most cases they will not even notice an interruption unless they are downloading large files. Users downloading a large file may have to restart their download after a failover.

Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.

Improving session synchronization performance

Two HA configuration options are available to reduce the performance impact of enabling session-pickup. They include reducing the number of sessions that are synchronized by adding a session pickup delay and using more FortiGate interfaces for session synchronization.

Reducing the number of sessions that are synchronized

If session pickup is enabled, as soon as new sessions are added to the primary unit session table they are synchronized to the other cluster units. Enable the `session-pickup-delay` CLI option to reduce the number of sessions that are synchronized by synchronizing sessions only if they remain active for more than 30 seconds.

Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
    set session-pickup-delay enable
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

Using multiple FortiGate interfaces for session synchronization

Using the `session-sync-dev` option you can select one or more FortiGate interfaces to use for synchronizing sessions as required for session pickup. Normally session synchronization occurs over the HA heartbeat link. Using this HA option means only the selected interfaces are used for session synchronization and not the HA heartbeat link. If you select more than one interface, session synchronization traffic is load balanced among the selected interfaces.

Moving session synchronization from the HA heartbeat interface reduces the bandwidth required for HA heartbeat traffic and may improve the efficiency and performance of the cluster, especially if the cluster is synchronizing a large number of sessions. Load balancing session synchronization among multiple interfaces can further improve performance and efficiency if the cluster is synchronizing a large number of sessions.

Use the following command to perform cluster session synchronization using the port10 and port12 interfaces.

```
config system ha
    set session-sync-dev port10 port12
end
```

Session synchronization packets use Ethertype 0x8892. The interfaces to use for session synchronization must be connected together either directly using the appropriate cable (possible if there are only two units in the cluster) or using switches. If one of the interfaces becomes disconnected the cluster uses the remaining interfaces for session synchronization. If all of the session synchronization interfaces become disconnected, session synchronization reverts back to using the HA heartbeat link. All session synchronization traffic is between the primary unit and each subordinate unit.

Since large amounts of session synchronization traffic can increase network congestion, it is recommended that you keep this traffic off of your network by using dedicated connections for it.

Session failover limitations for sessions passing through the cluster

This section contains information about session failover for communication sessions passing through the cluster. In general, if session pickup is enabled, session failover is supported for most TCP traffic. This section describes details about how this all works.

Protocol	Session Failover?
Most TCP sessions.	Supported if session-pickup is enabled. (More about TCP session failover on page 265)

Protocol	Session Failover?
IPv6, NAT64, and NAT66	Supported if session-pickup is enabled.
Proxy-based security profile sessions	Not Supported, sessions have to be restarted. Proxy-based features require the FortiGate to maintain very large amounts of internal state information for each session. The FGCP does not synchronize this internal state information. As a result, proxy-based sessions are not failed over. Active-active clusters can resume some of these sessions after a failover. (Active-active HA subordinate units sessions can resume after a failover on page 267)
Flow-based security profile sessions.	Supported if session-pickup is enabled. Flow-based sessions failover, but internal state information is not synchronized so sessions that fail over are no longer inspected by security profile functions. If both flow-based and proxy-based security profile features are applied to a TCP session, that session will not resume after a failover.
UDP and ICMP, multicast, or broadcast sessions	Supported if connectionless session-pickup is enabled. Otherwise, sessions have to be restarted. (UDP, ICMP, multicast and broadcast packet session failover on page 266)
GPRS Tunneling Protocol (GTP)	Supported with limitations. (FortiOS Carrier GTP session failover on page 266)
SIP	Supported for active-passive HA only. (SIP session failover on page 266)
SIMPLE, or SCCP signal session	Not supported, sessions have to be restarted.
SSL offloading and HTTP multiplexing	Not supported, sessions have to be restarted. (SSL offloading and HTTP multiplexing session failover on page 267)

More about TCP session failover

TCP sessions that are not being processed by security profile features resume after a failover even if these sessions are accepted by security policies with security profiles. Only TCP sessions that are actually being processed by these security profile features do not resume after a failover. For example:

- TCP sessions that are not virus scanned, web filtered, spam filtered, content archived, or are not SIP, SIMPLE, or SCCP signal traffic resume after a failover, even if they are accepted by a security policy with security profile options enabled. For example, SNMP TCP sessions through the FortiGate resume after a failover because FortiOS does not apply any security profile options to SNMP sessions.

- TCP sessions for a protocol for which security profile features have not been enabled resume after a failover even if they are accepted by a security policy with security profile features enabled. For example, if you have not enabled any antivirus or content archiving settings for FTP, FTP sessions resume after a failover.

UDP, ICMP, multicast and broadcast packet session failover

By default, even with session pickup enabled, the FGCP does not maintain a session table for UDP, ICMP, multicast, or broadcast packets. So the cluster does not specifically support failover of these packets.

Some UDP traffic can continue to flow through the cluster after a failover. This can happen if, after the failover, a UDP packet that is part of an already established communication stream matches a security policy. Then a new session will be created and traffic will flow. So after a short interruption, UDP sessions can appear to have failed over. However, this may not be reliable for the following reasons:

- UDP packets in the direction of the security policy must be received before reply packets can be accepted. For example, if a port1 -> port2 policy accepts UDP packets, UDP packets received at port2 destined for the network connected to port1 will not be accepted until the policy accepts UDP packets at port1 that are destined for the network connected to port2. So, if a user connects from an internal network to the Internet and starts receiving UDP packets from the Internet (for example streaming media), after a failover the user will not receive any more UDP packets until the user re-connects to the Internet site.
- UDP sessions accepted by NAT policies will not resume after a failover because NAT will usually give the new session a different source port. So only traffic for UDP protocols that can handle the source port changing during a session will continue to flow.

You can however, enable session pickup for UDP and ICMP packets by enabling session pickup for TCP sessions and then enabling session pickup for connectionless sessions:

```
config system ha
  set session-pickup enable
  set session-pickup-connectionless enable
end
```

This configuration causes the cluster units to synchronize UDP and ICMP session tables and if a failover occurs UDP and ICMP sessions are maintained.

SIP session failover

If session pickup is enabled, the FGCP supports SIP session failover (also called stateful failover) for active-passive HA.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in-progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

FortiOS Carrier GTP session failover

FortiOS Carrier HA supports GTP session failover. The primary unit synchronizes the GTP tunnel state to all cluster units after the GTP tunnel setup is completed. After the tunnel setup is completed, GTP sessions use UDP and HA does not synchronize UDP sessions to all cluster units. However, similar to other UDP sessions,

after a failover, since the new primary unit will have the GTP tunnel state information, GTP UDP sessions using the same tunnel can continue to flow with some limitations.

The limitation on packets continuing to flow is that there has to be a security policy to accept the packets. For example, if the FortiOS Carrier unit has an internal to external security policy, GTP UDP sessions using an established tunnel that are received by the internal interface are accepted by the security policy and can continue to flow. However, GTP UDP packets for an established tunnel that are received at the external interface cannot flow until packets from the same tunnel are received at the internal interface.

If you have bi-directional policies that accept GTP UDP sessions then traffic in either direction that uses an established tunnel can continue to flow after a failover without interruption.

SSL offloading and HTTP multiplexing session failover

SSL offloading and HTTP multiplexing are both enabled from firewall virtual IPs and firewall load balancing. Similar to the features applied by security profile, SSL offloading and HTTP multiplexing requires the FortiGate to maintain very large amounts of internal state information for each session. Sessions accepted by security policies containing virtual IPs or virtual servers with SSL offloading or HTTP multiplexing enabled do not resume after a failover.

Active-active HA subordinate units sessions can resume after a failover

In an active-active cluster, subordinate units process sessions. After a failover, all cluster units that are still operating may be able to continue processing the sessions that they were processing before the failover. These sessions are maintained because after the failover the new primary unit uses the HA session table to continue to send session packets to the cluster units that were processing the sessions before the failover. Cluster units maintain their own information about the sessions that they are processing and this information is not affected by the failover. In this way, the cluster units that are still operating can continue processing their own sessions without loss of data.

The cluster keeps processing as many sessions as it can. But some sessions can be lost. Depending on what caused the failover, sessions can be lost in the following ways:

- A cluster unit fails (the primary unit or a subordinate unit). All sessions that were being processed by that cluster unit are lost.
- A link failure occurs. All sessions that were being processed through the network interface that failed are lost.

This mechanism for continuing sessions is not the same as session failover because:

- Only the sessions that can be maintained are maintained.
- The sessions are maintained on the same cluster units and not re-distributed.
- Sessions that cannot be maintained are lost.

Session failover limitations for sessions terminated by the cluster

This section contains information about session failover for communication sessions terminated by the cluster. Sessions terminated by the cluster include management sessions as well as IPsec and SSL VPN, WAN Optimization and so on between the cluster and a client.

In general, most sessions terminated by the cluster have to be restarted after a failover. There are some exceptions though, for example, the FGCP provides failover for IPsec and SSL VPN sessions terminated by the cluster.



The session pickup setting does not affect session failover for sessions terminated by the cluster. Also other cluster settings such as active-active or active-passive mode do not affect session failover for sessions terminated by the cluster.

Protocol	Session Failover?
Administrative or management connections such as connecting to the GUI or CLI, SNMP, syslog, communication with FortiManager, FortiAnalyzer and so on.	Not supported, sessions have to be restarted.
Explicit web proxy, WCCP, WAN Optimization and Web Caching.	Not supported, sessions have to be restarted. (Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and Web Caching session failover on page 268)
IPsec VPN tunnels terminating at the FortiGate	Supported. SAs and related IPsec VPN tunnel data is synchronized to cluster members. (Synchronizing IPsec VPN SAs on page 269)
SSL VPN tunnels terminating at the FortiGate	Partially supported. Sessions are not synchronized and have to be restarted. Authentication failover and cookie failover is supported. Once the client restarts the session they shouldn't have to re-authenticate. (SSL VPN session failover and SSL VPN authentication failover on page 269)
PPTP and L2TP VPN terminating at the FortiGate	Not supported, sessions have to be restarted.

Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and Web Caching session failover

Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and web caching sessions all require the FortiGate to maintain very large amounts of internal state information for each session. This information is not maintained and these sessions do not resume after a failover.

The active-passive HA clustering is recommended for WAN optimization. All WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

Also, Web cache and byte cache databases are only stored on the primary unit. These databases are not synchronized to the cluster. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGates that it is participating with in WAN optimization tunnels.

SSL VPN session failover and SSL VPN authentication failover

Session failover is not supported for SSL VPN tunnels. However, authentication failover is supported for the communication between the SSL VPN client and the FortiGate. This means that after a failover, SSL VPN clients can re-establish the SSL VPN session between the SSL VPN client and the FortiGate without having to authenticate again.

However, all sessions inside the SSL VPN tunnel that were running before the failover are stopped and have to be restarted. For example, file transfers that were in progress would have to be restarted. As well, any communication sessions with resources behind the FortiGate that are started by an SSL VPN session have to be restarted.

To support SSL VPN cookie failover, when an SSL VPN session starts, the FGCP distributes the cookie created to identify the SSL VPN session to all cluster units.

PPTP and L2TP VPN sessions

PPTP and L2TP VPNs are supported in HA mode. For a cluster you can configure PPTP and L2TP settings and you can also add security policies to allow PPTP and L2TP pass through. However, the FGCP does not provide session failover for PPTP or L2TP. After a failover, all active PPTP and L2TP sessions are lost and must be restarted.

Synchronizing IPsec VPN SAs

The FGCP synchronizes IPsec security associations (SAs) between cluster members so that if a failover occurs, the cluster can resume IPsec sessions without having to establish new SAs. The result is improved failover performance because IPsec sessions are not interrupted to establish new SAs. Also, establishing a large number of SAs can reduce cluster performance.

The FGCP implements slightly different synchronization mechanisms for IKEv1 and IKEv2.

Synchronizing SAs for IKEv1

When an SA is synchronized to the subordinate units, the sequence number is set to the maximum sequence number. After a failover, all inbound traffic that connects with the new primary unit and uses the SA will be accepted without needing to re-key. However, first outbound packet to use the SA causes the sequence number to overflow and so causes the new primary unit to re-key the SA.

Please note the following:

- The cluster synchronizes all IPsec SAs.
- IPsec SAs are not synchronized until the IKE process has finished synchronizing the ISAKMP SAs. This is required in for dialup tunnels since it is the synchronizing of the ISAKMP SA that creates the dialup tunnel.
- A dialup interface is created as soon as the phase 1 is complete. This ensures that when HA synchronizes phase 1 information the dialup name is included.
- If the IKE process re-starts for any reason it deletes any dialup tunnels that exist. This forces the peer to re-key them.
- IPsec SA deletion happens immediately. Routes associated with a dialup tunnel that is being deleted are cleaned up synchronously as part of the delete, rather than waiting for the SA hard-expiry.

- The FGCP does not sync the IPsec tunnel MTU from the primary unit to the subordinate units. This means that after HA failover if the first packet received by the FortiGate arrives after the HA route has been deleted and before the new route is added and the packet is larger than the default MTU of 1024 then the FortiGate sends back an ICMP fragmentation required. However, as soon as routing is re-established then the MTU will be corrected and traffic will flow.

Synchronizing SAs for IKEv2

Due to the way the IKEv2 protocol is designed the FGCP cannot use exactly the same solution that is used for synchronizing IKEv1 SAs, though it is similar.

For IKEv2, like IKEv1, the FGCP synchronizes IKE and ISAKMP SAs from the primary unit to the subordinate units. However, for IKEv2 the FGCP cannot actually use this IKE SA to send/receive IKE traffic because IKEv2 includes a sequence number in every IKE message and thus it would require synchronizing every message to the subordinate units to keep the sequence numbers on the subordinate units up to date.

Instead, the FGCP synchronizes IKEv2 Message IDs. This Message ID Sync allows IKEv2 to re-negotiate send and receive message ID counters after a failover. By doing this, the established IKE SA can remain up, instead of re-negotiating.

The `diagnose vpn ike stats` command shows statistics for the number of HA messages sent/received for IKEv2. The output of this command includes a number of fields prefixed with `ha` that contain high availability related-data. For example:

```
.
.
.
ha.resync: 0
ha.vike.sync: 0
ha.conn.sync: 0
ha.sync.tx: 1
ha.sync.rx: 0
ha.sync.rx.len.bad: 0
.
.
.
```

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended HA configuration for WAN optimization is active-passive mode. Also, when the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions. HA also does not support WAN optimization session failover.

In a cluster, the primary unit only stores web cache and byte cache databases. These databases are not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGates that it is participating with in WAN optimization tunnels.

HA and load balancing

FGCP active-active (a-a) load balancing distributes network traffic among all of the units in a cluster. Load balancing can improve cluster performance because the processing load is shared among multiple cluster units.

This chapter describes how active-active load balancing works and provides detailed active-active HA NAT/Route and Transparent mode packet flow descriptions.

Load balancing overview

FGCP active-active HA uses a technique similar to unicast load balancing in which the primary unit is associated with the cluster HA virtual MAC addresses and cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit then uses a load balancing schedule to distribute sessions to all of the units in the cluster (including the primary unit). Subordinate unit interfaces retain their actual MAC addresses and the primary unit communicates with the subordinate units using these MAC addresses. Packets exiting the subordinate units proceed directly to their destination and do not pass through the primary unit first.

By default, active-active HA load balancing distributes proxy-based security profile processing to all cluster units. Proxy-based security profile processing is CPU and memory-intensive, so FGCP load balancing may result in higher throughput because resource-intensive processing is distributed among all cluster units.

Proxy-based security profile processing that is load balanced includes proxy-based virus scanning, proxy-based web filtering, proxy-based email filtering, and proxy-based data leak prevention (DLP) of HTTP, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, and NNTP, sessions accepted by security policies.

Other features enabled in security policies such as Endpoint security, traffic shaping and authentication have no effect on active-active load balancing.

You can also enable `load-balance-all` to have the primary unit load balance all TCP sessions. Load balancing TCP sessions increases overhead and may actually reduce performance so it is disabled by default. You can also enable `load-balance-udp` to have the primary unit load balance all UDP sessions. Load balancing UDP sessions also increases overhead so it is also disabled by default.

NP4 and NP6 processors can also offload and accelerate load balancing.

During active-active HA load balancing the primary unit uses the configured load balancing schedule to determine the cluster unit that will process a session. The primary unit stores the load balancing information for each load balanced session in the cluster load balancing session table. Using the information in this table, the primary unit can then forward all of the remaining packets in each session to the appropriate cluster unit. The load balancing session table is synchronized among all cluster units.

HTTPS, ICMP, multicast, and broadcast sessions are never load balanced and are always processed by the primary unit. IPS, Application Control, flow-based virus scanning, flow-based web filtering, flow-based DLP, flow-based email filtering, VoIP, IM, P2P, IPsec VPN, HTTPS, SSL VPN, HTTP multiplexing, SSL offloading, WAN optimization, explicit web proxy, and WCCP sessions are also always processed only by the primary unit.

In addition to load balancing, active-active HA also provides the same session, device and link failover protection as active-passive HA. If the primary unit fails, a subordinate unit becomes the primary unit and resumes operating the cluster.

Active-active HA also maintains as many load balanced sessions as possible after a failover by continuing to process the load balanced sessions that were being processed by the cluster units that are still operating. See [Session failover \(session-pickup\) on page 262](#) for more information.

Load balancing schedules

The load balancing schedule controls how the primary unit distributes packets to all cluster units. You can select from the following load balancing schedules.

Schedule	Description
None	No load balancing. Select None when the cluster interfaces are connected to load balancing switches. If you select None, the Primary unit does not load balance traffic and the subordinate units process incoming traffic that does not come from the Primary unit. For all other load balancing schedules, all traffic is received first by the Primary unit, and then forwarded to the subordinate units. The subordinate units only receive and process packets sent from the primary unit.
Hub	Load balancing if the cluster interfaces are connected to a hub. Traffic is distributed to cluster units based on the source IP and destination IP of the packet.
Least-Connection	If the cluster units are connected using switches, select Least Connection to distribute network traffic to the cluster unit currently processing the fewest connections.
Round-Robin	If the cluster units are connected using switches, select Round-Robin to distribute network traffic to the next available cluster unit.
Weighted Round-Robin	Similar to round robin, but weighted values are assigned to each of the units in a cluster based on their capacity and on how many connections they are currently processing. For example, the primary unit should have a lower weighted value because it handles scheduling and forwards traffic. Weighted round robin distributes traffic more evenly because units that are not processing traffic will be more likely to receive new connections than units that are very busy.
Random	If the cluster units are connected using switches, select Random to randomly distribute traffic to cluster units.
IP	Load balancing according to IP address. If the cluster units are connected using switches, select IP to distribute traffic to units in a cluster based on the source IP and destination IP of the packet.
IP Port	Load balancing according to IP address and port. If the cluster units are connected using switches, select IP Port to distribute traffic to units in a cluster based on the source IP, source port, destination IP, and destination port of the packet.

Once a packet has been propagated to a subordinate unit, all packets are part of that same communication session are also propagated to that same subordinate unit. Traffic is distributed according to communication session, not just according to individual packet.

Any subordinate unit that receives a forwarded packet processes it, without applying load balancing. Note that subordinate units are still considered to be active, because they perform routing, virus scanning, and other FortiGate tasks on their share of the traffic. Active subordinate units also share their session and link status information with all cluster units. The only things that active members do not do is make load balancing decisions.

Even though the primary unit is responsible for the load balancing process, the primary unit still acts like a FortiGate in that it processes packets, performing, routing, firewall, virus scanning, and other FortiGate tasks on its share of the traffic. Depending on the load balancing schedule used, the primary unit may assign itself a smaller share of the total load.

More about active-active failover

If a subordinate unit fails, the primary unit re-distributes the sessions that the subordinate was processing among the remaining active cluster members. If the primary unit fails, the subordinate units negotiate to select a new primary unit. The new primary unit continues to distribute packets among the remaining active cluster units.

Failover works in a similar way if the cluster consists of only two units. If the primary unit fails the subordinate unit negotiates and becomes the new primary unit. If the subordinate unit fails, the primary unit processes all traffic. In both cases, the single remaining unit continues to function as a primary unit, maintaining the HA virtual MAC address for all of its interfaces.

HTTPS sessions, active-active load balancing, and proxy servers

To prevent HTTPS web filtering problems active-active HA does not load balance HTTPS sessions. The FortiGate identifies HTTPS sessions as all sessions received on the HTTPS TCP port. The default HTTPS port is 443. You can go to **Policy & Objects > Policy > SSL/SSH Inspection** to use a custom port for HTTPS sessions. If you change the HTTPS port, the FGCP stops load balancing all sessions that use the custom HTTPS port.

Normally you would not change the HTTPS port. However, if your network uses a proxy server for HTTPS traffic you may have to change to the custom HTTPS port used by your proxy server. If your network uses a proxy server you might also use the same port for both HTTP and HTTPS traffic. In this case you would configure the FortiGate to use custom ports for both HTTP and HTTPS traffic. Go to **Policy & Objects > Policy > Proxy Options** to use a custom port for HTTP.

Using the same port for HTTP and HTTPS traffic can cause problems with active-active clusters because active-active clusters always load balance HTTP traffic. If both HTTP and HTTPS use the same port, the active-active cluster cannot differentiate between HTTP and HTTPS traffic and will load balance both.

As mentioned above, load balancing HTTPS traffic may cause problems with HTTPS web filtering. To avoid this problem, you should configure your proxy server to use different ports for HTTP and HTTPS traffic. Then configure your cluster to also use different ports for HTTP and HTTPS.

Selecting a load balancing schedule

You can select the load balancing schedule when initially configuring the cluster and you can change the load balancing schedule at any time while the cluster is operating without affecting cluster operation.

You can select a load balancing schedule from the CLI. Use the following command to select a load balancing schedule:

```
config system ha
    set schedule {hub | ip | ipport | leastconnection | none | random | round-robin
        | weight-round-robin}
end
```

Load balancing TCP and UDP sessions

You can use the following command to configure the cluster to load balance TCP sessions in addition to security profile sessions.

```
config system ha
    set load-balance-all enable
end
```

Enabling `load-balance-all` to add load balancing of TCP sessions may not improve performance because the cluster requires additional overhead to load balance sessions. Load balancing a TCP session usually requires about as much overhead as just processing it. On the other hand, TCP load balancing performance may be improved if your FortiGate includes NP4 or NP6 processors.

You can enable `load-balance-all` and monitor network performance to see if it improves. If performance is not improved, you might want to change the HA mode to active-passive since active-active HA is not providing any benefit.

On some FortiGate models you can use the following command to also load balance UDP sessions:

```
config system ha
    set load-balance-udp enable
end
```

Similar to load balancing TCP sessions, load balancing UDP sessions may also not improve performance. Also UDP load balancing performance may be improved with NP4 and NP6 processors.

Using NP4 or NP6 processors to offload load balancing

FortiGates that include NP4 and NP6 network processors can provide hardware acceleration for active-active HA cluster by offloading load balancing from the primary unit CPU. Network processors are especially useful when load balancing TCP and UDP sessions.

The first packet of every new session is received by the primary unit and the primary unit uses its load balancing schedule to select the cluster unit that will process the new session. This information is passed back to the network processor and all subsequent packets of the same sessions are offloaded to the network processor which sends the packet directly to a subordinate unit. Load balancing is effectively offloaded from the primary unit to the network processor resulting in a faster and more stable active-active cluster.

To take advantage of network processor load balancing acceleration, connect the cluster unit interfaces with network processors to the busiest networks. Connect non-accelerated interfaces to less busy networks. No special FortiOS or HA configuration is required. Network processor acceleration of active-active HA load balancing is supported for any active-active HA configuration or active-active HA load balancing schedule.

Configuring weighted-round-robin weights

You can configure weighted round-robin load balancing for a cluster and configure the static weights for each of the cluster units according to their priority in the cluster. When you set `schedule` to `weight-round-robin` you can use the `weight` option to set the static weight of each cluster unit. The static weight is set according to the priority of each unit in the cluster. A FortiGate HA cluster can contain up to four FortiGates so you can set up to 4 static weights.

The priority of a cluster unit is determined by its device priority, the number of monitored interfaces that are functioning, its age in the cluster and its serial number. Priorities are used to select a primary unit and to set the priorities of all of the subordinate units. Thus the priority of a cluster unit can change depending on configuration settings, link failures and so on. Since weights are also set using this priority, the weights are independent of specific cluster units but do depend on the role of the each unit in the cluster.

You can use the following command to display the relative priorities of the units in a cluster. The cluster unit serial numbers and their priorities are listed in the last few lines of the command output. This example shows a cluster of three FortiGates:

```
get system ha status
.
.
.
Slave :1 FG-5KD3914800284
Master:0 FG-5KD3914800344
Slave :2 FG-5KD3914800353
```

The primary unit always has the highest priority and the subordinate units have lower priorities.

The default static weight for each cluster unit is 40. This means that sessions are distributed evenly among all cluster units. You can use the `set weight` command to change the static weights of cluster units to distribute sessions to cluster units depending on their priority in the cluster. The weight can be between 0 and 255. Increase the weight to increase the number of connections processed by the cluster unit with that priority.

You set the weight for each unit separately. For the example cluster of 3 FortiGates you can set the weight for each unit as follows:

```
config system ha
    set mode a-a
    set schedule weight-round-robin
    set weight 0 5
    set weight 1 10
    set weight 2 15
end
```

If you enter the `get` command to view the HA configuration the output for `weight` would be:

```
weight 5 10 15 40 40 40 40 40 40 40 40 40 40 40
```

This configuration has the following results if the output of the `get system ha status` command is that shown above:

- The first five connections are processed by the primary unit (priority 0, weight 5).
- The next 10 connections are processed by the first subordinate unit (priority 1, weight 10)
- The next 15 connections are processed by the second subordinate unit (priority 2, weight 15)

Dynamically optimizing weighted load balancing according to how busy cluster units are

In conjunction with using static weights to load balance sessions among cluster units you can configure a cluster to dynamically load balance sessions according to individual cluster unit CPU usage, memory usage, and number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy-based security profile sessions. If any of these system loading indicators increases above configured thresholds, weighted load balancing dynamically sends fewer new sessions to the busy unit until it recovers.

High CPU or memory usage indicates that a unit is under increased load and may not be able to process more sessions. HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy use are also good indicators of how busy a cluster unit is, since processing high numbers of these proxy sessions can quickly reduce overall cluster unit performance.

For example, you can set a CPU usage high watermark threshold. When a cluster unit reaches this high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to the low watermark threshold. When the low watermark threshold is reached the cluster resumes normal load balancing of sessions to the cluster unit.

You can set individual high and low watermark thresholds and weights for CPU usage, memory usage, and for the number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions.

The CPU usage, memory usage, and proxy weights determine how the cluster load balances sessions when a high watermark threshold is reached and also affect how the cluster load balances sessions when multiple cluster units reach different high watermark thresholds at the same time. For example, you might be less concerned about a cluster unit reaching the memory usage high watermark threshold than reaching the CPU usage high watermark threshold. If this is the case you can set the weight lower for memory usage. Then, if one cluster unit reaches the CPU usage high watermark threshold and a second cluster unit reaches the memory usage high watermark threshold the cluster will load balance more sessions to the cluster unit with high memory usage and fewer sessions to the cluster unit with high CPU usage. As a result, reaching the CPU usage high watermark will have a greater affect on how sessions are redistributed than reaching the memory usage high watermark.

When a high watermark threshold is reached, the corresponding weight is subtracted from the static weight of the cluster unit. The lower the weight the fewer the number of sessions that are load balanced to that unit.

Subsequently when the low watermark threshold is reached, the static weight of the cluster unit returns to its configured value. For the weights to all be effective the weights assigned to the load indicators should usually be lower than or equal to the static weights assigned to the cluster units.

Use the following command to set thresholds and weights for CPU and memory usage and HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold <weight> <low> <high>
  set memory-threshold <weight> <low> <high>
  set http-proxy-threshold <weight> <low> <high>
  set ftp-proxy-threshold <weight> <low> <high>
  set imap-proxy-threshold <weight> <low> <high>
  set nntp-proxy-threshold <weight> <low> <high>
  set pop3-proxy-threshold <weight> <low> <high>
  set smtp-proxy-threshold <weight> <low> <high>
end
```

For each option, the weight range is 0 to 255 and the default weight is 5. The low and high watermarks are a percent (0 to 100). The default low and high watermarks are 0 which means they are disabled. The default configuration when weighted load balancing is enabled looks like the following:

```
config system ha
    set mode a-a
    set schedule weight-round-robin
    set cpu-threshold 5 0 0
    set memory-threshold 5 0 0
    set http-proxy-threshold 5 0 0
    set ftp-proxy-threshold 5 0 0
    set imap-proxy-threshold 5 0 0
    set nntp-proxy-threshold 5 0 0
    set pop3-proxy-threshold 5 0 0
    set smtp-proxy-threshold 5 0 0
end
```



When you first enable HA weighted load balancing, the weighted load balancing configuration is synchronized to all cluster units and each cluster unit has the default configuration shown above. Changes to the CPU, memory, HTTP, FTP, IMAP, NNTP, POP3, and SMTP proxy thresholds and low and high watermarks must be made for each cluster unit and are not synchronized to the other cluster units.

When you configure them, the high watermarks must be greater than their corresponding low watermarks.

For CPU and memory usage the low and high watermarks are compared with the percentage CPU and memory use of the cluster unit. For each of the proxies the high and low watermarks are compared to a number that represents percent of the max number of proxy sessions being used by a proxy. This number is calculated using the formula:

$$\text{proxy usage} = (\text{current sessions} * 100) / \text{max sessions}$$

where:

`current sessions` is the number of active sessions for the proxy type.

`max sessions` is the session limit for the proxy type. The session limit depends on the FortiGate and its configuration.

You can use the following command to display the maximum and current number of sessions for a proxy:

```
get test {ftpd | http | imap | nntp | pop3 | smtp} 4
```

You can use the following command to display the maximum number of sessions and the and current number of sessions for all of the proxies:

```
get test proxyworker 4
```

The command output includes lines similar to the following:

```
get test http 4
HTTP Common
Current Connections          5000/8032
```

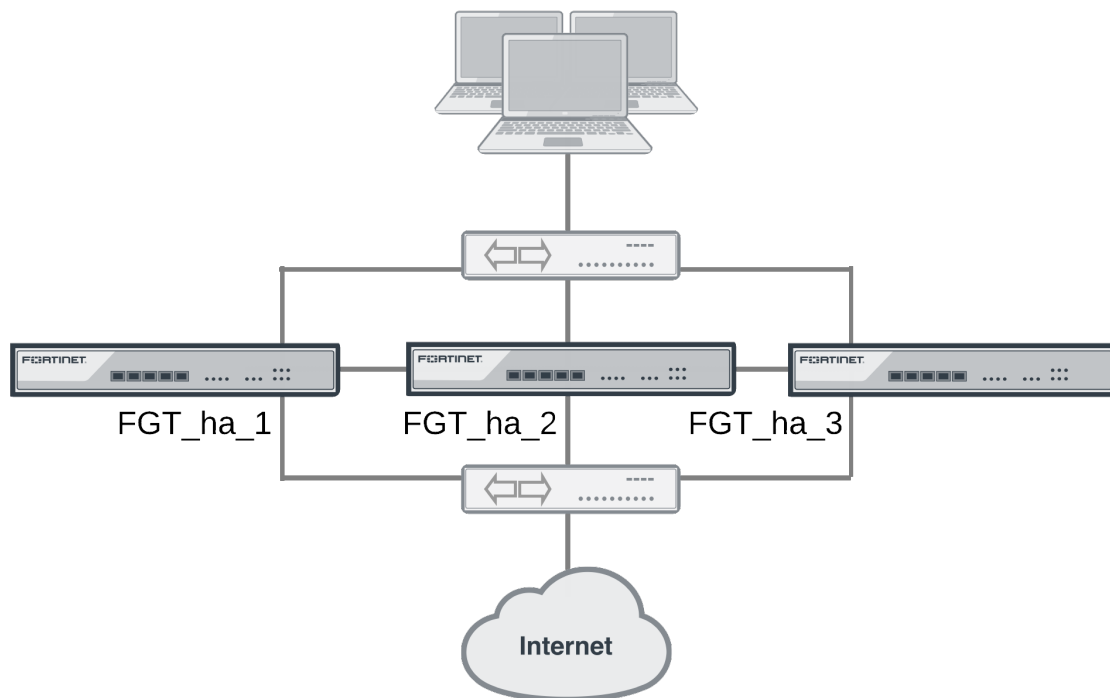
In the example, 5000 is the current number of proxy connections being used by HTTP and 8032 is the maximum number of proxy sessions allowed. For this example the proxy usage would be:

```
proxy usage = (5000 * 100) / 8032
proxy usage = 62%
```

Example weighted load balancing configuration

Consider a cluster of three FortiGate with host names FGT_ha_1, FGT_ha_2, and FGT_ha_3 as shown below. This example describes how to configure weighted load balancing settings for CPU and memory usage for the cluster and then to configure HTTP and POP3 proxy weights to send most HTTP and POP3 proxy sessions to different cluster units.

Example HA weighted load balancing configuration



Connect to the cluster CLI and use the following command to set the CPU usage threshold weight to 30, low watermark to 60, and high watermark to 80. This command also sets the memory usage threshold weight to 10, low watermark to 60, and high watermark to 90.

```

config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold 30 60 80
  set memory-threshold 10 60 90
end
  
```

The static weights for the cluster units remain at the default values of 40. Since this command changes the mode to `a-a` and the schedule to `weight-round-robin` for the first time, the weight settings are synchronized to all cluster units.

As a result of this configuration, if the CPU usage of any cluster unit (for example, FGT_ha_1) reaches 80% the static weight for that cluster unit is reduced from 40 to 10 and only 10 of every 120 new sessions are load

balanced to this cluster unit. If the memory usage of FGT_ha_1 also reaches 90% the static weight further reduces to 0 and no new sessions are load balanced to FGT_ha_1. Also, if the memory usage of FGT_ha_2 reaches 90% the static weight of FGT_ha_2 reduces to 30 and 30 of every 120 new sessions are load balanced to FGT_ha_2.

Now that you have established the weight load balancing configuration for the entire cluster you can monitor the cluster to verify that processing gets distributed evenly to all cluster units. From the GUI you can go to **System > HA > View HA Statistics** and see the CPU usage, active sessions, memory usage and other statistics for all of the cluster units. If you notice that one cluster unit is more or less busy than others you can adjust the dynamic weights separately for each cluster unit.

For example, in some active-active clusters the primary unit may tend to be busier than other cluster units because in addition to processing sessions the primary unit also receives all packets sent to the cluster and performs load balancing to distribute the sessions to other cluster units. To reduce the load on the primary unit you could reduce the CPU and memory usage high watermark thresholds for the primary unit so that fewer sessions are distributed to the primary unit. You could also reduce the primary unit's high watermark setting for the proxies to distribute more proxy sessions to other cluster units.



This would only be useful if you are using device priorities and override settings to make sure the same unit always becomes the primary unit. See [Controlling primary unit selection using device priority and override on page 49](#).

If the example cluster is configured for FGT_ha_2 to be the primary unit, connect to the FGT_ha_2's CLI and enter the following command to set CPU usage, memory usage, and proxy usage high watermark thresholds lower.

```
config system ha
  set cpu-threshold 30 60 70
  set memory-threshold 30 60 70
  set http-proxy-threshold 30 60 70
  set ftp-proxy-threshold 30 60 70
  set imap-proxy-threshold 30 60 70
  set nntp-proxy-threshold 30 60 70
  set pop3-proxy-threshold 30 60 70
  set smtp-proxy-threshold 30 60 70
end
```

As a result, when any of these factors reaches 70% on the primary unit, fewer sessions will be processed by the primary unit, preventing the number of sessions being processed from rising.

NAT/Route mode active-active cluster packet flow

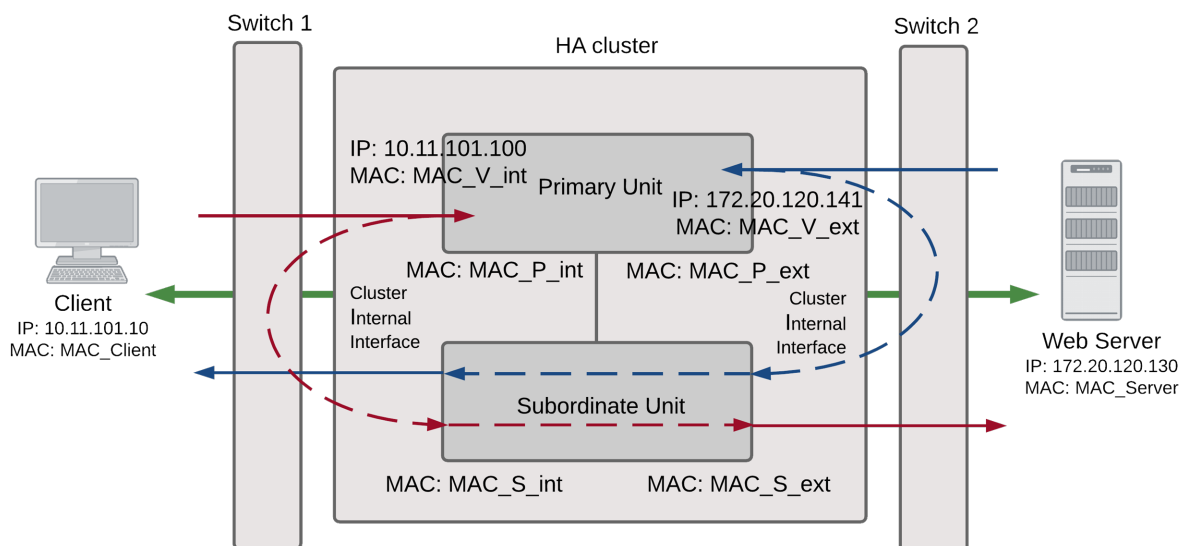
This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In NAT/Route mode, eight MAC addresses are involved in active-active communication between the client and the web server when the primary unit load balances packets to the subordinate unit:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only knows the cluster external virtual MAC address (MAC_V_ext).

NAT/Route mode active-active packet flow



Packet flow from client to web server

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
5. The client's request packet reaches the primary unit internal interface.

IP address	MAC address
------------	-------------

Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_P_int
Destination	172.20.120.130	MAC_S_int

- The subordinate unit recognizes that the packet has been forwarded from the primary unit and processes it.
- The subordinate unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_S_ext
Destination	172.20.120.130	MAC_Server

- The primary unit forwards further packets in the same session to the subordinate unit.
- Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from web server to client

- When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
- The web server issues an ARP request to 172.20.120.141.
- The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.
- The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
--	-------------------	--------------------

Source	172.20.120.130	MAC_P_ext
Destination	172.20.120.141	MAC_S_ext

6. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
7. The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_S_int
Destination	10.11.101.10	MAC_Client

8. The primary unit forwards further packets in the same session to the subordinate unit.
9. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails, the subordinate unit negotiates to become the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
3. The new primary unit sends gratuitous ARP packets to the 10.10.101.0 network to associate its internal IP address with the internal virtual MAC address.
4. The new primary unit sends gratuitous ARP packets to the 172.20.120.0 network to associate its external IP address with the external virtual MAC address.
5. Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

Transparent mode active-active cluster packet flow

This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.

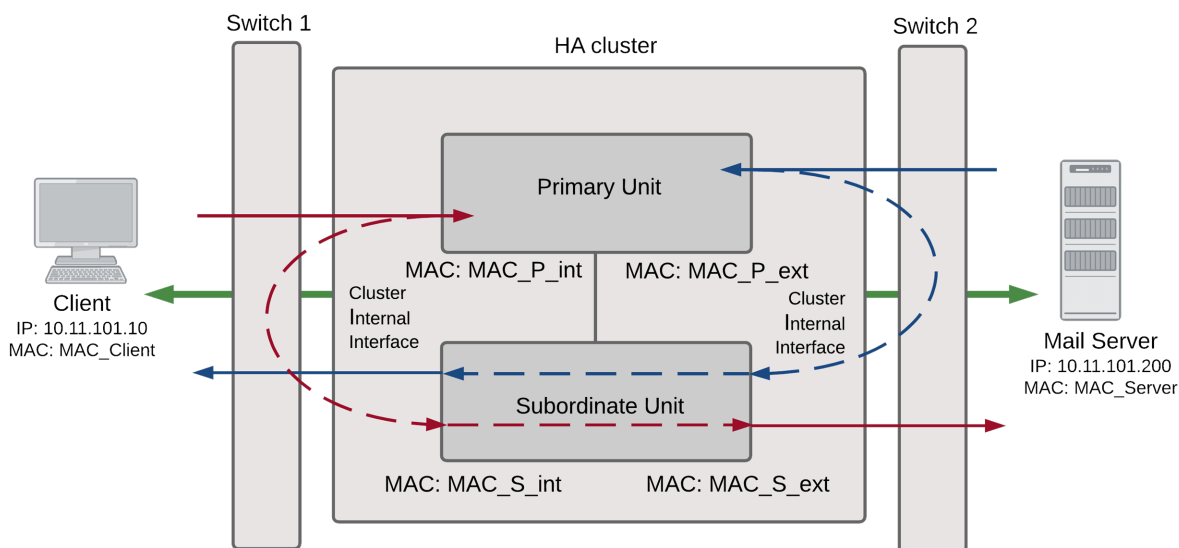
In Transparent mode, six MAC addresses are involved in active-active communication between a client and a server when the primary unit load balances packets to the subordinate unit:

- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

The HA virtual MAC addresses are not directly involved in communicate between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and load balanced among cluster members.

The cluster's presence on the network and its load balancing are transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the external virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the internal virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Transparent mode active-active packet flow



Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

6. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_P_int
Destination	10.11.101.200	MAC_S_int

7. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
8. The subordinate unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_S_ext
Destination	10.11.101.200	MAC_Server

9. The primary unit forwards further packets in the same session to the subordinate unit.
10. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from mail server to client

1. To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
2. The primary unit forwards the ARP request to the client computer.
3. The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
4. The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

5. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_P_ext
Destination	10.11.101.10	MAC_S_ext

6. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
7. The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_S_int
Destination	10.11.101.10	MAC_Client

8. The primary unit forwards further packets in the same session to the subordinate unit.
9. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails the subordinate unit negotiates to become the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
3. The new primary unit sends gratuitous ARP requests to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
4. The new primary unit sends gratuitous ARP requests to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
5. Traffic sent to the cluster is now received and processed by the new primary unit.

If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

HA with FortiGate-VM and third-party products

This chapter provides information about operating FortiOS VM cluster and operating FortiGate clusters with third party products such as layer-2 and layer-3 switches.

FortiGate-VM for VMware HA configuration

If you want to combine two or more FortiGate-VM instances into a FortiGate Clustering Protocol (FGSP) High Availability (HA) cluster the VMware server's virtual switches used to connect the heartbeat interfaces must operate in promiscuous mode. This permits HA heartbeat communication between the heartbeat interfaces. HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The FGCP uses link-local IPv4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

To enable promiscuous mode in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect heartbeat interfaces.
4. In the **Properties** window left pane, select **vSwitch** and then select **Edit**.
5. Select the **Security** tab, set **Promiscuous Mode** to **Accept**, then select **OK**.
6. Select **Close**.

You must also set the virtual switches connected to other FortiGate interfaces to allow MAC address changes and to accept forged transmits. This is required because the FGCP sets virtual MAC addresses for all FortiGate interfaces and the same interfaces on the different VM instances in the cluster will have the same virtual MAC addresses.

To make the required changes in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect FortiGate VM interfaces.
4. Set **MAC Address Changes** to **Accept**.
5. Set **Forged Transmits** to **Accept**.

FortiGate VM for Hyper-V HA configuration

Promiscuous mode and support for MAC address spoofing is required for FortiGate-VM for Hyper-V to support FortiGate Clustering Protocol (FGCP) high availability (HA). By default the FortiGate-VM for Hyper-V has promiscuous mode enabled in the XML configuration file in the FortiGate-VM Hyper-V image. If you have problems with HA mode, confirm that this is still enabled.

In addition, because the FGCP applies virtual MAC addresses to FortiGate data interfaces and because these virtual MAC addresses mean that matching interfaces of different FortiGate-VM instances will have the same virtual MAC addresses you have to configure Hyper-V to allow MAC spoofing. But you should only enable MAC spoofing for FortiGate-VM data interfaces. You should not enable MAC spoofing for FortiGate HA heartbeat interfaces.

With promiscuous mode enabled and the correct MAC spoofing settings you should be able to configure HA between two or more FortiGate-VM for Hyper-V instances.

Troubleshooting layer-2 switches

Issues may occur because of the way an HA cluster assigns MAC addresses to the primary unit. Two clusters with the same group ID cannot connect to the same switch and cannot be installed on the same network unless they are separated by a router.

Forwarding delay on layer 2 switches

You must ensure that if there is a switch between the FortiGate HA cluster and the network it is protecting and the switch has a forwarding delay (even if spanning tree is disabled) when one of its interfaces is activated then the forwarding delay should be set as low as possible. For example, some versions of Cisco IOS have a forwarding delay of 15 seconds even when spanning tree is disabled. If left at this default value then TCP session pickup can fail because traffic is not forwarded through the switch on HA failover.

Failover issues with layer-3 switches

After a failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables of the switches connected to the cluster. If the cluster is connected using layer-2 switches, the MAC forwarding tables (also called arp tables) are refreshed by the gratuitous ARP packets and the switches start directing packets to the new primary unit.

In some configurations that use layer-3 switches, after a failover, the layer-3 switches may not successfully re-direct traffic to the new primary unit. The possible reason for this is that the layer-3 switch might keep a table of IP addresses and interfaces and may not update this table for a relatively long time after the failover (the table is not updated by the gratuitous ARP packets). Until the table is updated, the layer-3 switch keeps forwarding packets to the now failed cluster unit. As a result, traffic stops and the cluster does not function.

As of the release date of this document, Fortinet has not developed a workaround for this problem. One possible solution would be to clear the forwarding table on the layer-3 switch.

The `config system ha link-failed-signal` command described in [Updating MAC forwarding tables when a link failover occurs on page 247](#) can be used to resolve link failover issues similar to those described here.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Ethertype conflicts with third-party switches

Some third-party network equipment may use packets with Ethertypes that are the same as the ethertypes used for HA heartbeat packets. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 heartbeat packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGates connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

You can also use the following CLI commands to change the Ethertypes of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For more information, see [Heartbeat packet Ethertypes on page 221](#).

LACP, 802.3ad aggregation and third-party switches

If a cluster contains 802.3ad aggregated interfaces you should connect the cluster to switches that support configuring multiple Link Aggregation (LAG) groups.

The primary and subordinate unit interfaces have the same MAC address, so if you cannot configure multiple LAG groups a switch may place all interfaces with the same MAC address into the same LAG group; disrupting the operation of the cluster.

You can change the FortiGate configuration to prevent subordinate units from participating in LACP negotiation. For example, use the following command to do this for an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```

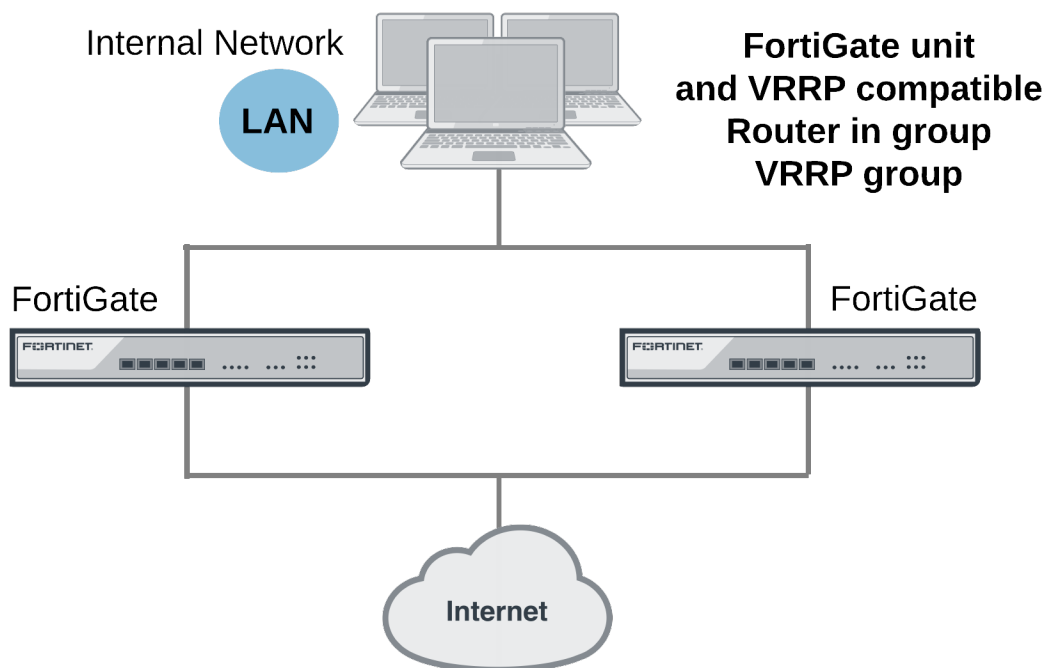

This configuration prevents the subordinate unit interfaces from sending or receiving packets. Resulting in the cluster not being able to operate in active-active mode. As well, failover may be slower because after a failover the new primary unit has to perform LACP negotiation before being able to process network traffic.

For more information, see [FGCP HA with 802.3ad aggregated interfaces on page 115](#).

VRRP

A Virtual Router Redundancy Protocol (VRRP) configuration can be used as a high availability solution to make sure that a network maintains connectivity with the Internet (or with other networks) even if the default router for the network fails. Using VRRP, if a router or a FortiGate fails all traffic to this router transparently fails over to another router or FortiGate that takes over the role of the router or FortiGate that failed. If the failed router or FortiGate is restored, it will once again take over processing traffic for the network. VRRP is described by [RFC 3768](#).

Example VRRP configuration



To configure VRRP you create a VRRP group that contains two or more routers. Some or all of these routers can be FortiGates. You can include different FortiGate models in the same VRRP group. The group members are configured to be the master router and one or more backup routers of the VRRP group. The network directs all traffic to the master's IP address and MAC address. If the master fails, VRRP dynamically shifts packet forwarding to a backup router. VRRP provides this redundancy without user intervention or additional configuration to any of the devices on the network.

The VRRP redundancy scheme means that devices on the network keep a single IP address for the default gateway and this IP address maps to a well-known virtual MAC address. If the VRRP master fails, one of the backup units becomes the new master and acquires virtual IP and MAC addresses that match the addresses of the master. The network then automatically directs all traffic to the backup unit. VRRP uses the broadcast capabilities of Ethernet networks. As long as one of the routers in a VRRP group is running, ARP requests for the

default gateway IP address always receive replies. Additionally, hosts can send packets outside their subnet without interruption.

FortiGates support VRRP and can be quickly and easily integrated into a network that has already deployed a group of routers using VRRP. You can also create a new VRRP configuration consisting of a FortiGate acting as a VRRP master with one or more VRRP-compatible routers acting as backup routers. Some or all of those backup routers can be FortiGates.

During normal operation the VRRP master unit sends VRRP advertisement messages to the backup units. A backup unit will not attempt to become a master unit while it is receiving these messages. When a FortiGate operating as a VRRP master fails, a backup unit takes its place and continues processing network traffic. The backup unit assumes the master unit has failed if it stops receiving the advertisement messages from the master unit. The backup unit with the highest priority becomes the new master unit after a short delay. During this delay the new master unit sends gratuitous ARPs to the network to map the virtual router IP address to its MAC address. As a result, all packets sent to the default route IP address are sent to the new master unit. If the backup unit is a FortiGate, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate is back online.

During a VRRP failover, as the backup unit starts to forward traffic it will not have session information for all of the failed over in-progress sessions. If the backup unit is operating as a normal FortiGate it will not be able to forward this traffic because of the lack of session information. To resolve this problem, immediately after a failover and for a short time as its taking over traffic processing, the backup unit operates with asymmetric routing enabled. This allows the backup unit to re-create all of the in-progress sessions and add them to the session table. While operating with asymmetric routing enabled, the backup unit cannot apply security functions. When the start-time ends the backup unit disables asymmetric routing and returns to normal operation including applying security functions.

Adding a VRRP virtual router to a FortiGate interface

Use the following command to add a VRRP virtual router to the port10 interface of a FortiGate. This VRRP virtual router has a virtual router ID of 200, uses IP address 10.31.101.200 and has a priority of 255. Since this is the highest priority this interface is configured to be the master of the VRRP group with ID number 200.



VRRP can be configured only on physical interfaces or VLAN interfaces. You cannot configure VRRP on hardware-switch interfaces where multiple physical interfaces are combined into a hardware switch interface.

```
config system interface
edit port10
config vrrp
edit 200
set vrip 10.31.101.200
set priority 255
end
end
```

VRRP virtual MAC address

The VRRP virtual MAC address (or virtual router MAC address) is a shared MAC address adopted by the VRRP master. If the master fails the same virtual MAC master fails over to the new master. As a result, all packets for VRRP routers can continue to use the same virtual MAC address. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

Each VRRP router is associated with its own virtual MAC address. The last part of the virtual MAC depends on the VRRP virtual router ID using the following format:

```
00-00-5E-00-01-<VRID_hex>
```

Where <VRID_hex> is the VRRP virtual router ID in hexadecimal format in internet standard bit-order. For more information about the format of the virtual MAC see [RFC 3768](#).

Some examples:

- If the VRRP virtual router ID is 10 the virtual MAC would be 00-00-5E-00-01-0a.
- If the VRRP virtual router ID is 200 the virtual MAC would be 00-00-5E-00-01-c8.

The VRRP virtual MAC address feature is disabled by default. When you enable the feature on a FortiGate interface, all of the VRRP routers added to that interface use their own VRRP virtual MAC address. Each virtual MAC address will be different because each virtual router has its own ID.

Use the following command to enable the VRRP virtual MAC address on the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  end
end
```

The port2 interface will now accept packets sent to the MAC addresses of the VRRP virtual routers added to this interface.

Using the VRRP virtual MAC address can improve network efficiency especially on large and complex LANs because when a failover occurs devices on the LAN do not have to learn a new MAC address for the new VRRP router.

If the VRRP virtual MAC address feature is disabled, the VRRP group uses the MAC address of the master. In the case of a FortiGate VRRP virtual router this is the MAC address of the FortiGate interface that the VRRP virtual routers are added to. If a master fails, when the new master takes over it sends gratuitous ARPs to associate the VRRP virtual router IP address with the MAC address of the new master (or the interface of the FortiGate that has become the new master). If the VRRP virtual MAC address is enabled the new master uses the same MAC address as the old master.

VRRP Groups

A VRRP group includes all the relevant VRRP IDs and tracks the VRRP status to force the status of all group members if a VRRP domain is changed from master to backup. VRRP groups are configured through the CLI. The VRRP group ID can be between 1 and 65535. Use the following command to add a VRRP group to the port20 interface that includes the virtual route identifiers 25, 50, 66 and 70 to VRRP group 10

```
config system interface
  edit port20
```

```
config vrrp
  edit 25
    set vrgrp 10
  next
  edit 50
    set vrgrp 10
  next
  edit 66
    set vrgrp 10v
  next
  edit 70
    set vrgrp 10
end
```

Using a Second Destination IP (VRDST)

VRRP can now be configured with second destination IP (VRDST) for monitoring. When two IPs are used, VRRP failure will only be reported if both monitored IPs are down.

Use the following command to configure a second destination IP (VRDST) to port14:

```
config system interface
  edit port14
    config vrrp
      edit 12
        set vrdst 10.10.10.20 10.20.20.10
      end
    end
```

Configuring VRRP

To configure VRRP you must configure two or more FortiGate interfaces or routers with the same virtual router ID and IP address. Then these FortiGates or routers can automatically join the same VRRP group. You must also assign priorities to each of the FortiGates or routers in the VRRP group. One of the FortiGates or routers must have the highest priority to become the master. The other FortiGates or routers in the group are assigned lower priorities and become backup units. All of the units in the VRRP group should have different priorities. If the master unit fails, VRRP automatically fails over to the remaining unit in the group with the highest priority.

You configure VRRP from the FortiGate CLI by adding a VRRP virtual router to a FortiGate interface. You can add VRRP virtual routers to multiple FortiGate interfaces and you can add more than one virtual router to the same interface.

Example VRRP configuration: two FortiGates in a VRRP group

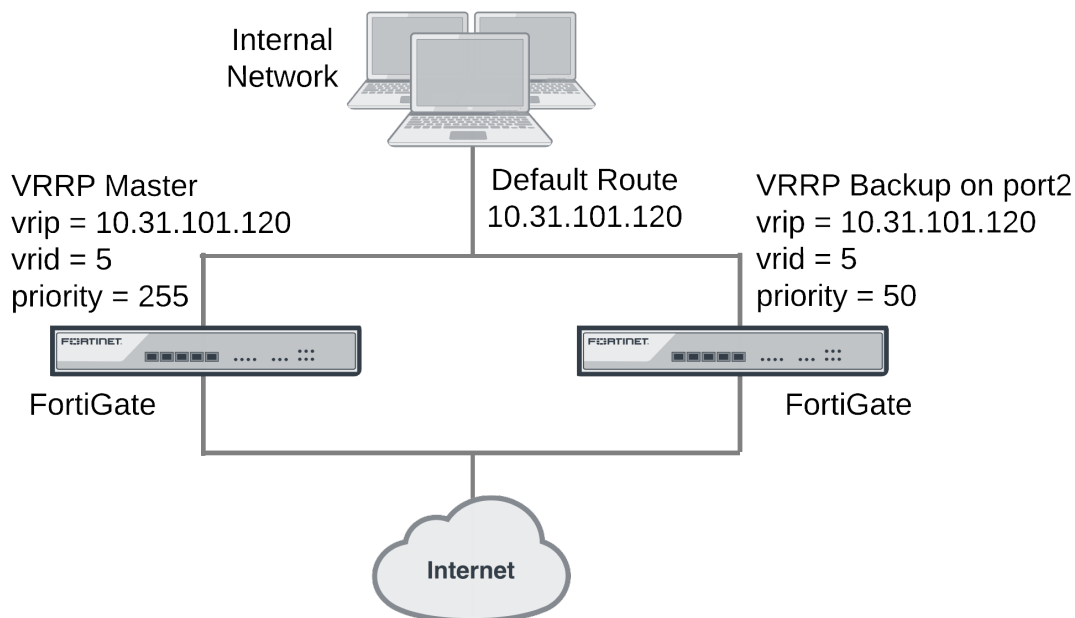
This example includes a VRRP group consisting of two FortiGates that connect an internal network to the Internet. As shown below, the internal network's default route is 10.31.101.120.

The FortiGate port2 interfaces connect to the internal network. A VRRP virtual router is added to each FortiGate's port2 interface. The virtual router IP address is 10.31.101.120 (the internal network's default route) and the virtual router's ID is 5. The VRRP priority of the master unit is set to 255 and the VRRP priority of the backup unit is 50.

The port2 interface of each FortiGate should have an IP address that is different from the virtual router IP address and the port2 interface IP addresses should be different from each other.

This example also includes enabling the VRRP virtual MAC address on both FortiGate port2 interfaces so that the VRRP group uses the VRRP virtual MAC address.

Example VRRP configuration with two FortiGates



To configure the FortiGates for VRRP

1. Select one of the FortiGates to be the VRRP master and the other to be the backup unit.
2. From the master unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  config vrrp
    edit 5
      set vrip 10.31.101.120
      set priority 255
    end
  end
```

3. From the backup unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  config vrrp
    edit 5
```

```
set vrip 10.31.101.120
set priority 50
end
end
```

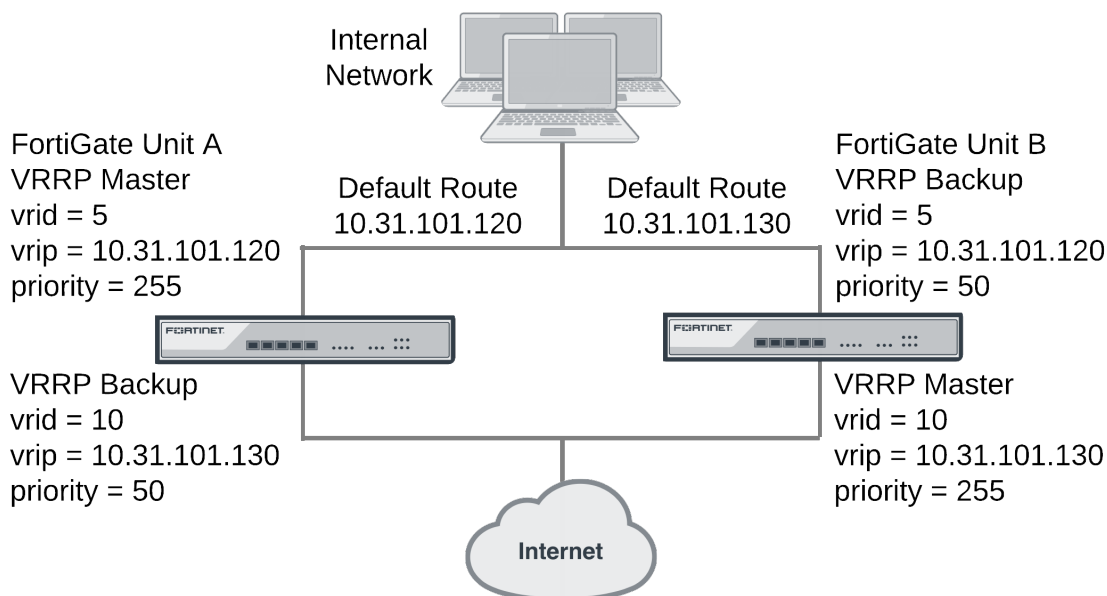
Example VRRP configuration: VRRP load balancing two FortiGates and two VRRP groups

In this configuration two VRRP groups are involved. Each FortiGate participates in both of them. One FortiGate is the master of one group and the other FortiGate is the master of the other group. The network distributes traffic between two different default routes (10.31.101.120 and 10.31.101.130). One VRRP group is configured with one of the default route IP addresses and the other VRRP group get the other default route IP address. So during normal operation both FortiGates are processing traffic and the VRRP groups are used to load balance the traffic between the two FortiGates.

If one of the FortiGates fails, the remaining FortiGate becomes the master of both VRRP groups. The network sends all traffic for both default routes to this FortiGate. The result is a configuration that under normal operation load balances traffic between two FortiGates, but if one of the FortiGates fails, all traffic fails over to the unit that is still operating.

This example also includes enabling the VRRP virtual MAC address on both FortiGate port2 interfaces so that the VRRP groups use their VRRP virtual MAC addresses.

Example VRRP configuration with two FortiGates and two VRRP groups



To configure the FortiGates

1. Log into the CLI of FortiGate A.
2. Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate A:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  config vrrp
    edit 50 (32)
      set vrip 10.31.101.120
      set priority 255
    next
    edit 100 (64)
      set vrip 10.31.101.130
      set priority 50
    end
  end
```

3. Log into the CLI of FortiGate B.
4. Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate B:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  config vrrp
    edit 50
      set vrip 10.31.101.120
      set priority 50
    next
    edit 100
      set vrip 10.31.101.130
      set priority 255
    end
  end
```

Optional VRRP configuration settings

In addition to the basic configuration settings, you can change to the VRRP configuration to:

- Adjust the virtual router advertisement message interval between 1 and 255 seconds using the `adv-interval` option.
- Adjust the startup time using the `start-time` option. The default start time is 3 seconds and the range is 1 to 255 seconds. The start time is the maximum time that the backup unit waits between receiving advertisement messages from the master unit. If the backup unit does not receive an advertisement message during this time it assumes the master has failed and becomes the new master unit. In some cases the advertisement messages may be delayed. For example, some switches with spanning tree enabled may delay some of the advertisement message packets. If you find that backup units are attempting to become master units without the master unit failing, you can extend the start time to make sure the backup units wait long enough for the advertisement messages.

- Enable or disable individual virtual router configurations using the `status` option. Normally virtual router configurations are enabled but you can temporarily disable one if its not required.
- Enable or disable preempt mode using the `preempt` option. In preempt mode a higher priority backup unit can preempt a lower priority master unit. This can happen if a master has failed, a backup unit has become the master unit, and the failed master is restarted. Since the restarted unit will have a higher priority, if preempt mode is enabled the restarted unit will replace the current master unit. Preempt mode is enabled by default.
- Monitor the route to a destination IP address using the `vrdst` option.

FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two or more FortiGates can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute sessions among the FortiGates and the FGSP performs session synchronization of IPv4 and IPv6 TCP, UDP, ICMP, expectation, and NAT sessions and IPsec tunnels to keep the session tables of the FortiGates synchronized. If one of the FortiGates fails, session failover occurs and active sessions fail over to the FortiGates that are still operating. This failover occurs without any loss of data. As well, the external routers or load balancers will detect the failover and re-distribute all sessions to the peers that are still operating.

The FortiGates operate as peers that process traffic and synchronize sessions with other FortiGates in the cluster. An FGSP cluster can include from 2 to 16 FortiGates. Adding more FortiGates increases the CPU and memory required to keep all of the FortiGates synchronized. So depending on your network conditions, adding too many FortiGates to an FGSP cluster may reduce overall performance.

The FortiGates in the FGSP cluster must be the same model and be running the same firmware version. You use the `config system cluster-sync` command to configure FGSP between the FortiGates and the `config system ha` command to configure what is synchronized.

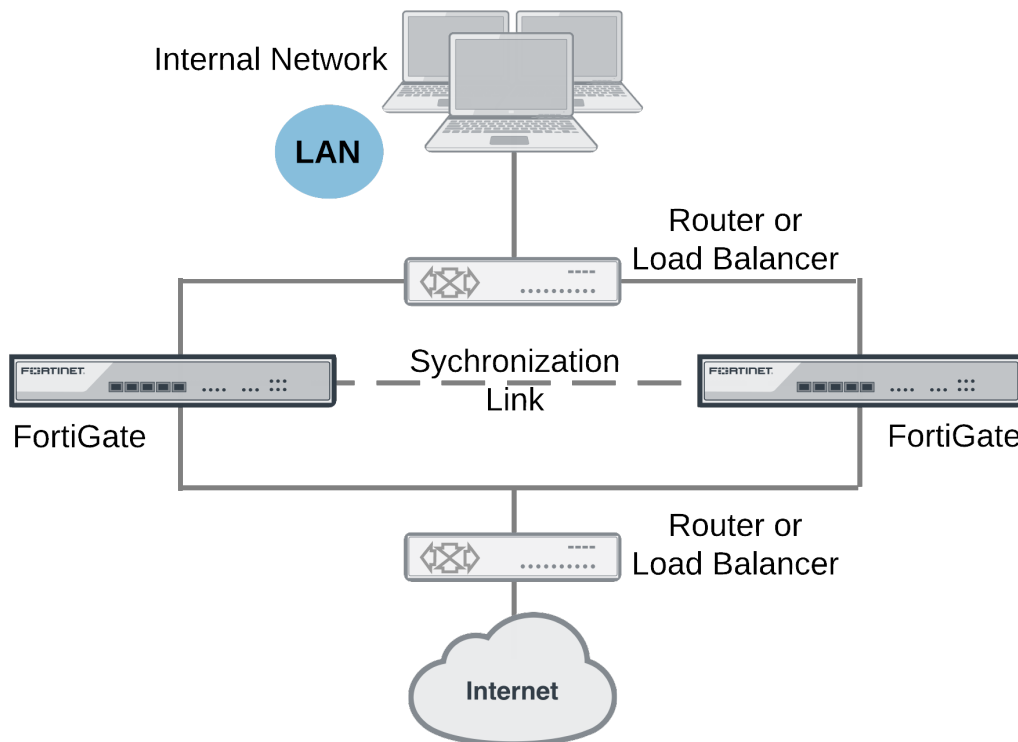


In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. The FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions and IPsec tunnels.



You cannot configure FGSP when FGCP is enabled. However FGSP is compatible with VRRP.

The FGSP can be used instead of FGCP HA to provide **session synchronization** between two peer FortiGates. If the external load balancers direct all sessions to one peer the affect is similar to active-passive FGCP HA. If external load balancers or routers load balance traffic to both peers, the effect is similar to active-active FGCP HA. The load balancers should be configured so that all of the packets for any given session are processed by the same peer. This includes return packets.

FGSP HA

By default, FGSP synchronizes all IPv4 and IPv6 TCP sessions, IPsec tunnels, and also synchronizes the configuration of the FortiGates.

You can optionally enable session pickup to synchronize connectionless (UDP and ICMP) sessions, expectation sessions, and NAT sessions. If you do not enable session pickup, the FGSP does not share session tables for the particular session type and sessions do not resume after a failover. All sessions that are interrupted by the failover and must be re-established at the application level. Many protocols can successfully restart sessions with little, or no, loss of data. Others may not recover easily. Enable session pickup for sessions that may be difficult to reestablish. Since session pickup requires FortiGate memory and CPU resources, only enable this feature for sessions that you need to have synchronized.

The synchronization link is set up in the same way as FGCP heartbeat interfaces. You must connect the synchronization link interfaces together and use the heartbeat device (hbdev) option to add the heartbeat devices to the configuration.

You can also optionally add filters to control which sessions are synchronized. You can add filters to only synchronize packets from specified source and destination addresses, specified source and destination interfaces, and specified services.

Load balancing and session failover is done by external routers or load balancers instead of by the FGSP. The FortiGates just perform session synchronization to support session failover.

Configuring FGSP HA cluster-sync instances

You use the following command to configure an FGSP HA cluster-sync instance.

```
config system cluster-sync
  edit 1
    set peerip <peer-ip-address>
    set peervd <vdom-name>
    set syncvd <vdom-name>
  end
```

Where:

`peerip` is the IP address of an interface of another FortiGate in the FGSP cluster that this configuration synchronizes sessions to.

`peervd` is the name of the VDOM on the other FortiGate that should be synchronized with this one. By default the `peervd` is `root`.

`syncvd` is the name of the VDOM of the FortiGate that should be synchronized with the other FortiGate. If multiple VDOMs are not enabled `syncvd` should be set to `root`.



For FGSP HA to work properly, all VDOMs to be synchronized must be added to all of the FortiGates in cluster. The names of the matching interfaces in each VDOM must also be the same; this includes the names of matching VLAN interfaces. Note that the index numbers of the matching interfaces and VLAN interfaces can be different. Also the VLAN IDs of the matching VLAN interfaces can be different. If you enable configuration synchronization this will happen automatically.

This command creates a cluster-sync instance that causes a FortiGate to synchronize the TCP sessions of one of its VDOMs (by default the root VDOM) to the root VDOM of another FortiGate (which would become another FortiGate in the FGSP cluster). You can also use the `config system ha` command to synchronize more session types and to synchronize the configuration. Cluster-sync instances are not synchronized and must be added to each FortiGate in the cluster.

A cluster of two FortiGates would only require one `cluster-sync` instance for each VDOM to be synchronized. This instance would synchronize the sessions from the root VDOM of one FortiGate to the root VDOM of the other. The second FortiGate would also include a cluster-sync instance to synchronize its root VDOM with the other FortiGate's root VDOM.

In a multiple VDOM configuration, you add a separate cluster-sync instance for each VDOM to be synchronized. You don't have to synchronize all VDOMs. If multiple VDOMs are enabled, the `config system cluster-sync` command is a global command.

FGSP clusters with three or more FortiGates

If an FGSP cluster includes three or more FortiGates you must explicitly define all of the cluster-sync instances that you need. In a cluster of four FortiGates, each FortiGate can be synchronized with up to three other FortiGates so to synchronize all of the FortiGates you must add three cluster-sync instances to each FortiGate (or $n-1$, where n is the number of FortiGates in the cluster).

Selecting the sessions to synchronize

You can add a cluster-sync instance with a filter to only synchronize some sessions. A filter can be added to a cluster-sync instance as follows:

```
config system cluster-sync
  edit 1
    set peerip <peer-ip-address>
    set peervd <vdom-name>
    set syncvd <vdom-name>
    config session-sync-filter
      srcintf <interface-name>
      dstintf <interface-name>
      srcaddr x.x.x.x x.x.x.x
      dstaddr x.x.x.x x.x.x.x
      srcaddr6 ::/x
      dstaddr6 ::/x
    end
  end
end
```

You can use the filter to only synchronize sessions according to the session source and destination interface and IPv4 or IPv6 address.

You can only add one filter to a cluster-sync instance. To create multiple filters you must create multiple cluster-sync instances.

Synchronizing sessions

Use the following command to enable session synchronization and configure the FGSP to use the port8 interface for synchronizing traffic:

```
config system ha
  set session-pickup enable
  set hbdev "port8" 50
end
```

Synchronizing the configuration

The FGSP includes configuration synchronization, allowing you to make configuration changes once for all of the FortiGates in the cluster instead of requiring you to make redundant configuration changes on each FortiGate. Settings that identify the FortiGate to the network, for example, interface IP addresses and BGP neighbor settings, and are not synchronized so each FortiGate maintains its identity on the network. As well cluster-sync instances are not synchronized and must be configured separately on each FortiGate in the cluster.

By default configuration synchronization is disabled. You can use the following command to enable it.

```
config system ha
  set standalone-config-sync enable
end
```

IPsec tunnel synchronization

When you use the `config system cluster-sync` command to enable FGSP, IPsec keys and other runtime data (but not actual tunnel sessions) are synchronized between cluster units. This means that if one of the cluster units goes down the cluster units that are still operating can quickly get IPsec tunnels re-established without re-negotiating them. However, after a failover all existing tunnel sessions on the failed FortiGate have to be restarted on the still operating FortiGates.

IPsec tunnel sync only supports dialup IPsec. The interfaces on the FortiGates that are tunnel endpoints must have the same IP addresses and external routers must be configured to load balance IPsec tunnel sessions to the FortiGates in the cluster.

Synchronizing UDP and ICMP (connectionless) sessions

In many configurations, due to their non-stateful nature, UDP and ICMP sessions don't need to be synchronized to naturally failover. However, if required you can configure the FGSP to synchronize UDP and ICMP sessions by entering the following:

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
end
```

Synchronizing NAT sessions

By default, NAT sessions are not synchronized. However, the FGSP can synchronize NAT sessions if you enter the following:

```
config system ha
    set session-pickup enable
    set session-pickup-nat enable
end
```

However, if you want NAT sessions to resume after a failover you should not configure NAT to use the destination interface IP address since the FGSP FortiGates have different IP addresses. With this configuration, after a failover all sessions that include the IP addresses of interfaces on the failed FortiGate will have nowhere to go since the IP addresses of the failed FortiGate will no longer be on the network.

Instead, in an FGSP configuration, if you want NAT sessions to failover you should use IP pools with the type set to overload (which is the default IP pool type). For example:

```
config firewall ippool
    edit FGSP-pool
        set type overload
        set startip 172.20.120.10
        set endip 172.20.120.20
    end
```

Then when you configure NAT firewall policies, turn on NAT and select to use dynamic IP pool and select the IP Pool that you added. Configuration synchronization should add the same IP pools and firewall policies to all

FortiGates in the cluster. If configuration synchronization is not enabled you must add the same IP pools and policies to all of the FortiGates in the cluster.

Synchronizing asymmetric sessions

By default, asymmetric sessions are not synchronized. Normally, session synchronization cannot be asymmetric because it is stateful. So all of the packets of a given session must be processed on the same FortiGate. This includes return packets.

However, if you have an asymmetric routing configuration, you can enter the following command to synchronize asymmetric sessions by dynamically detecting asymmetric sessions and disabling anti-reply for these sessions.

```
config system ha
  set session-pickup enable
  set session-pickup-expectation enable
end
```

The FGSP enforces firewall policies for asymmetric traffic, including cases where the TCP 3-way handshake is split between two FortiGates. For example, FGT-A receives the TCP-SYN, FGT-B receives the TCP-SYN-ACK, and FGT-A receives the TCP-ACK. Under normal conditions a firewall will drop this connection since the 3-way handshake was not seen by the same firewall. However two FortiGates with FGSP configured will be able to properly pass this traffic since the firewall sessions are synchronized.

This asymmetric function can also work with connectionless UDP and ICMP traffic. If traffic will be highly asymmetric, as described above, the following command must be enabled on both FortiGates.

```
config system ha
  set session-pickup enable
  set session-pickup-connectionless enable
end
```

Synchronizing asymmetric traffic can be very useful in situations where multiple Internet connections from different ISPs are spread across multiple FortiGates. Since it is typically not possible to guarantee Internet bound traffic leaving via an ISP will return using the exact same ISP, the FGSP provides critical firewall functions in this situation.



Asymmetric sessions may not be synchronized in low latency networks if the reply packet is received before the peer has received the session synchronization packet. This limitation usually only occurs in low latency networks.

The FGSP also has applications in virtualized computing environments where virtualized hosts move between data centers. The firewall session synchronization features of FGSP allow for more flexibility than in traditional firewalling functions.

Synchronizing expectation sessions

FortiOS session helpers keep track of the communication of Layer-7 protocols such as FTP and SIP that have control sessions and expectation sessions. Usually the control sessions establish the link between server and client and negotiate the ports and protocols that will be used for data communications. The session helpers then create expectation sessions through the FortiGate for the ports and protocols negotiated by the control session.

The expectation sessions are usually the sessions that actually communicate data. For FTP, the expectation sessions transmit files being uploaded or downloaded. For SIP, the expectation sessions transmit voice and video data. Expectation sessions usually have a timeout value of 30 seconds. If the communication from the server is not initiated within 30 seconds the expectation session times out and traffic will be denied.

By default the FGSP does not synchronize expectation sessions and if a failover occurs the sessions will have to be restarted.

If you want to synchronize expectation sessions so that they will continue after a failover you can enter the following:

```
config system ha
    set session-pickup enable
    set session-pickup-expectation enable
end
```

Security profile flow-based inspection and asymmetric traffic

Security profile inspection (flow or proxy based) for a session is not expected to work properly if the traffic in the session is balanced across more than one FortiGate in either direction. Flow-based inspection should be used in FGSP deployments.

For an environment where traffic is symmetric, security profile inspection can be used with the following limitations:

- No session synchronization for the sessions inspected using proxy-based inspection. Sessions will drop and need to be reestablished after data path failover.
- Sessions with flow-based inspection will failover; however, inspection of failed over sessions after the failover may not work.

A single FortiGate must see both the request and reply traffic for security profile inspection to function correctly. For environments where asymmetric traffic is expected, security profile inspection should not be used.

Notes and limitations

FGSP HA has the following limitations:

- The FGSP is a global configuration option. As a result you can only add one service to a filter configuration. You cannot add custom services or service groups even if virtual domains are not enabled.
- You can only add one filter configuration to a given FGSP configuration. However, you can add multiple filters by adding multiple identical FGSP configurations, each one with a different filter configuration.
- Sessions accepted by security policies with security profiles configured are not synchronized.
- FGSP HA is configured from the CLI.
- FGSP HA is available for FortiGates or virtual domains operating in NAT/Route or Transparent mode. NAT sessions are not synchronized in either mode (unless NAT synchronization is enabled as described in [Synchronizing NAT sessions on page 302](#)). In NAT/Route mode, only sessions for route mode security policies are synchronized. In Transparent mode, only sessions for normal Transparent mode policies are synchronized.

- FGSP HA is supported for traffic on physical interfaces, VLAN interfaces, zones, aggregate interfaces, and NPx (NP4, NP6 etc.) accelerated interfaces. The FGSP has not been tested for inter-vdom links, between HA clusters, and for redundant interfaces.
- The names of the matching interfaces, including VLAN interfaces, aggregate interfaces and so on, must be the same on both peers.
- An FGSP cluster can include from 2 to 16 FortiGates. Adding more FortiGates increases the CPU and memory required to keep all of the FortiGates synchronized.

Configuring session synchronization links

When FGSP HA is operating, the FortiGates share session information over Ethernet links similar to an HA heartbeat link. Usually you would use the same interface on each FortiGate for session synchronization. If possible you should connect the session synchronization interfaces directly without using a switch or other networking equipment. For FortiGate-5000 systems you can use a backplane interface as the session synchronization link.

You can use different interfaces on each FortiGate for session synchronization links. Also, if you have multiple session synchronization configurations, you can have multiple links between the FortiGates. In fact if you are synchronizing a lot of sessions, you may want to configure and connect multiple session synchronization links to distribute session synchronization traffic to these multiple links.

You cannot configure backup session synchronization links. Each configuration only includes one session synchronization link.

The session synchronization link should always be maintained. If session synchronization communication is interrupted and a failure occurs, sessions will not failover and data could be lost.

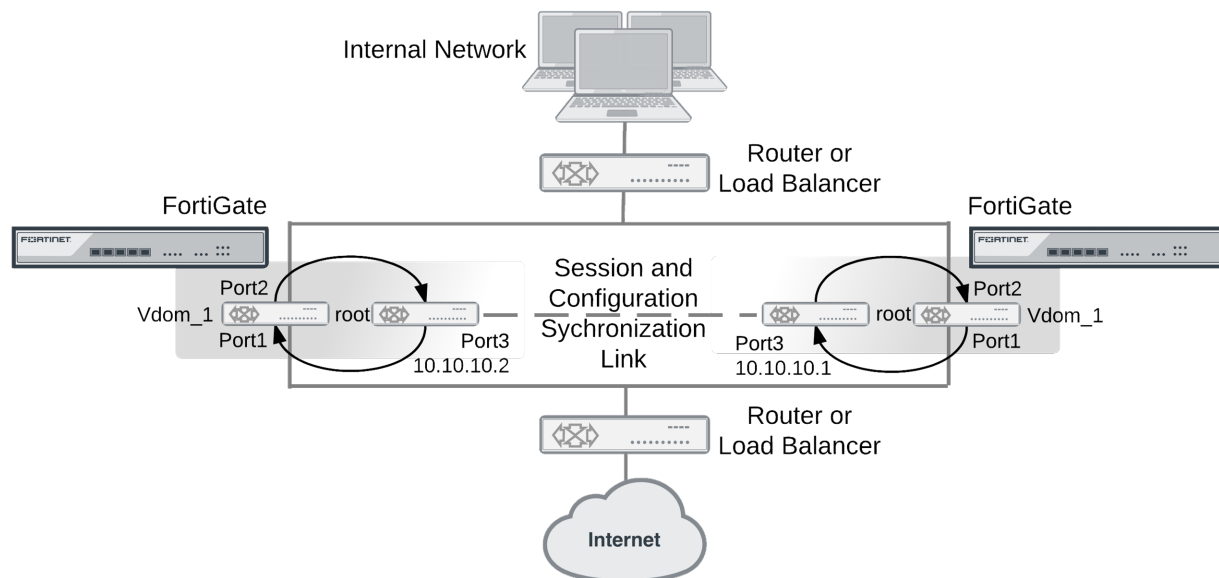
Session synchronization traffic can use a considerable amount of network bandwidth. If possible, session synchronization link interfaces should only be used for session synchronization traffic and not for data traffic.

Basic example configuration

The following configuration example shows how to configure basic FGSP HA for the two peer FortiGates shown below. The host names of peers are peer_1 and peer_2. Both peers are configured with two virtual domains: root and vdom_1. All sessions processed by vdom_1 are synchronized. The synchronization link interface is port3 which is in the root virtual domain. The IP address of port3 on peer_1 is 10.10.10.1. The IP address of port3 on peer_2 is 10.10.10.2.

Also on both peers, port1 and port2 are added to vdom_1. On peer_1 the IP address of port1 is set to 192.168.20.1 and the IP address of port2 is set to 172.110.20.1. On peer_2 the IP address of port1 is set to 192.168.20.2 and the IP address of port2 is set to 172.110.20.2.

Example FGSP HA network configuration



To configure FGSP HA

1. Configure the load balancer or router to send all sessions to peer_1.
2. Configure the load balancer or router to send all traffic to peer_2 if peer_1 fails.
3. Use normal FortiGate configuration steps on peer_1:
 - Enable virtual domain configuration.
 - Add the vdom_1 virtual domain.
 - Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
 - Set the IP address of port1 to 192.168.20.1.
 - Set the IP address of port2 to 172.110.20.1.
 - Set the IP address of port3 to 10.10.10.1.
 - Add route mode security policies between port1 and port2 to vdom_1.

4. Enter the following commands to configure session synchronization for peer_1:

```
config system cluster-sync
  edit 1
    set peerip 10.10.10.2
    set peervd root
    set syncvd vdom_1
  end
```

5. Use normal FortiGate configuration steps on peer_2:

- Enable virtual domain configuration.
 - Add the vdom_1 virtual domain.
 - Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
 - Set the IP address of port1 to 192.168.20.2.
 - Set the IP address of port2 to 172.110.20.2.
 - Set the IP address of port3 to 10.10.10.1.
 - Add route mode security policies between port1 and port2 to vdom_1.
6. Enter the following command to configure session synchronization for peer_1

```
config system cluster-sync
  edit 1
    set peerip 10.10.10.1
    set peervd root
    set syncvd vdom_1
  end
```

Now that the FortiGates are connected and configured their configurations are synchronized, so when you make a configuration change on one FortiGate it is synchronized to the other one.

To add filters

You can add a filter to this basic configuration if you only want to synchronize some TCP sessions. For example you can enter the following command to add a filter so that only HTTP sessions are synchronized:

```
config system cluster-sync
  edit 1
    config filter
      set service HTTP
    end
  end
```

You can also add a filter to control the source and destination addresses of the IPv4 packets that are synchronized. For example you can enter the following command to add a filter so that only sessions with source addresses in the range 10.10.10.100 to 10.10.10.200 are synchronized.

```
config system cluster-sync
  edit 1
    config filter
      set srcaddr 10.10.10.100 10.10.10.200
    end
  end
```

You can also add a filter to control the source and destination addresses of the IPv6 packets that are synchronized. For example you can enter the following command to add a filter so that only sessions with destination addresses in the range 2001:db8:0:2::/64 are synchronized.

```
config system cluster-sync
  edit 1
    config filter
      set dstaddr6 2001:db8:0:2::/64
    end
  end
```

To synchronize TCP sessions

You enter the following command to synchronization TCP sessions and set the synchronization link (heartbeat device):

```
config system ha
  set hbdev "port3" 50
  set session-pickup enable
end
```

To synchronize UDP and ICMP sessions

You enter the following command to add synchronization of UDP and ICMP sessions to this configuration:

```
config system ha
  set session-pickup enable
  set session-pickup-connectionless enable
end
```

To synchronize the configuration

Enter the following command to enable configuration synchronization.

```
config system ha
  set standalone-config-sync enable
end
```

Verifying FGSP configuration and synchronization

You can use the following diagnose commands to verify that the FGSP and its synchronization functions are operating correctly.

FGSP configuration summary and status

Enter the following command to display a summary of the FGSP configuration and synchronization status:

```
diagnose sys session sync
sync_ctx: sync_started=1, sync_tcp=1, sync_others=1,
sync_expectation=1, sync_redir=0, sync_nat=1, stdalone_sessync=0.
sync: create=12:0, update=0, delete=0:0, query=14
recv: create=14:0, update=0, delete=0:0, query=12
ses pkts: send=0, alloc_fail=0, recv=0, recv_err=0 sz_err=0
nCfgr_sess_sync_num=5, mtu=16000
sync_filter:
1: vd=0, szone=0, dzone=0, saddr=0.0.0.0:0.0.0.0, daddr=0.0.0.0:0.0.0.0,
```

`sync_started=1` shows that synchronization is working. If this is set to 0 then something is not correct with session synchronization and synchronization has not been able to start because of it.

`sync_tcp=1, sync_others=1, sync_expectation=1, and sync_nat=1` show that the FGSP has been configured to synchronize TCP, connectionless, asymmetric, and NAT sessions.

`sync: create=12:0` and `recv: create=14:0` show that this FortiGate has synchronized 12 sessions to its peer and has received 14 sessions from its peer.

`sync_filter` shows the configured FGSP filter. In this case no filter has been created so all sessions are synchronized.

`vd=0` indicates that root VDOM sessions are synchronized.

Verifying that sessions are synchronized

Enter the command `diagnose sys session list` to display information about the sessions being processed by the FortiGate. In the command output look for sessions that should be synchronized and make sure they contain output lines that include `syncd` for example, `state=log may_dirty ndr syncd`) to confirm that they are being synchronized by the FGSP.

```
diagnose sys session list
session info: proto=6 proto_state=05 duration=469 expire=0 timeout=3600
flags=00000000 sockflag=00000000 sockport=21 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=log may_dirty ndr syncd
statistic(bytes/packets/allow_err): org=544/9/1 reply=621/7/0 tuples=2
origin->sink: org pre->post, reply pre->post dev=46->45->46
gwy=10.2.2.1/10.1.1.1
hook=pre dir=org act=noop 192.168.1.50:45327->172.16.1.100:21(0.0.0.0:0)
hook=post dir=reply act=noop 172.16.1.100:21->192.168.1.50:45327(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=00002deb tos=ff/ff ips_view=1 app_list=2000 app=16427
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=192.168.1.50, bps=633
```



FORTINET®

High Performance Network Security



Copyright© 2017 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.